

PROTOSTAR: Generic Efficient Accumulation/Folding for Special-sound Protocols

Benedikt Bünz
Stanford University,
Espresso Systems

Binyi Chen
Espresso Systems

May 26, 2023

Abstract

Accumulation is a simple yet powerful primitive that enables incrementally verifiable computation (IVC) without the need for recursive SNARKs. We provide a generic, efficient accumulation (or folding) scheme for any $(2k - 1)$ -move special-sound protocol with a verifier that checks ℓ degree- d equations. The accumulation verifier only performs $k + 2$ elliptic curve multiplications and $k + d + O(1)$ field/hash operations. Using the compiler from BCLMS21 (Crypto 21), this enables building efficient IVC schemes where the recursive circuit only depends on the number of rounds and the verifier degree of the underlying special-sound protocol but not the proof size or the verifier time. We use our generic accumulation compiler to build PROTOSTAR. PROTOSTAR is a non-uniform IVC scheme for Plonk that supports high-degree gates and (vector) lookups. The recursive circuit is dominated by 3 group scalar multiplications and a hash of d^* field elements, where d^* is the degree of the highest gate. The scheme does not require a trusted setup or pairings, and the prover does not need to compute any FFTs. The prover in each accumulation/IVC step is also only logarithmic in the number of supported circuits and independent of the table size in the lookup.

Contents

1	Introduction	3
1.1	Technical overview	7
2	Preliminaries	11
2.1	Special-sound Protocols and Fiat-Shamir Transform	11
2.2	Adaptive Fiat-Shamir transform	12
2.3	Commitment Scheme	12
2.4	Incremental Verifiable Computation (IVC)	13
2.5	Simple Accumulation	14
3	Protocols	15
3.1	Special-sound Protocols	15
3.2	Commit and Open	16
3.3	Fiat-Shamir transform	17
3.4	Accumulation Scheme for V_{NARK}	18
3.5	Compressing verification checks for high-degree verifiers	24
4	Special-sound subprotocols for Protostar	29
4.1	Permutation relation	30
4.2	High-degree custom gate relation	30
4.3	Lookup relation	30
4.4	Vector-valued lookup	33
4.5	Circuit selection	36
5	Special-sound protocols for Plonkup relations	37
6	Protostar	39
A	Accumulation Scheme for high/low degree verifier	47
B	Protostar for CCS	49

1 Introduction

Incrementally Verifiable Computation [Val08] is a powerful primitive that enables a possibly infinite computation to be run, such that the correctness of the state of the computation can be verified at any point. IVC, and its generalization to DAGs, PCD [CT10], have many applications, including distributed computation [BCCT13; CTV15], blockchains [BMRS20; KB20], verifiable delay functions [BBBF18], verifiable photo editing [NT16], and SNARKs for machine-computations [BCTV14]. An IVC-based VDF construction is the current candidate VDF for Ethereum [KMT22]. One of the most exciting applications of IVC and PCD is the ZK-EVM. This is an effort to build a proof system that can prove that Ethereum blocks, as they exist today, are valid [But22].

Accumulation and folding. Historically, IVC was built from recursive SNARKs, proving that the previous computation step had a valid SNARK that proves correctness up to that point. Recently, an exciting new approach was initiated by Halo [BGH19] and has led to a series of significant advances [BCMS20; BCLMS21; KST22]. The idea is related to batch verification. Instead of verifying a SNARK at every step of the computation, we can instead *accumulate* the SNARK verification check with previous checks. We define an *accumulator*¹ such that we can combine a new SNARK and an old accumulator into a new accumulator. Checking or *deciding* the new accumulator implies that all previously accumulated SNARKs were valid. Now the recursive statement just needs to ensure the accumulation was performed correctly. Amazingly, this accumulation step can be significantly cheaper than SNARK verification [BGH19; BCMS20]. Even more surprising, this process does not even require a SNARK but instead can be instantiated with a non-succinct NARK [BCLMS21], as long as there exists an efficient accumulation scheme for that NARK. The most efficient accumulation (aka folding) scheme constructions yield IVC constructions, where the recursive circuit is dominated by as few as 2 elliptic curve scalar multiplications [BCLMS21; KST22]. These constructions only require the discrete logarithm assumption in the random oracle model and, unlike many efficient SNARK-based IVCs, do not require a trusted setup, pairings, or FFTs. These constructions build an accumulation scheme for one fixed (but universal) R1CS language by taking a random linear combination between the accumulator and a new proof. R1CS is a minimal expression of NP, defined by three matrices A, B, C , that closely resembles arithmetic circuits with addition and multiplication gates. However, it has limited flexibility, especially as the current constructions require fixing R1CS matrices that are used for all computation steps. These limitations are especially problematic for ZK-EVMs. In a ZK-EVM, each VM instruction (OP-CODE) is encoded in a different circuit. Each circuit uses high-degree gates, instead of just multiplication, and so-called lookup gates [GWC19]. These lookup gates enable looking up that a circuit value is in some table, simplifying range proofs and bit-operations.

¹Unrelated to set accumulators.

These R1CS-based accumulation schemes contrast non-IVC SNARK developments, with an increased focus on high-degree gate [GWC19; CBBZ22] and lookup support [GW20]. For lookups, a recent line of work has shown that if the table can be pre-computed, we can perform n lookups in a table of size T in time $O(n \log n)$, independent of T [ZBKMNS22; PK22; ZGKMR22; EFG22].

More expressive accumulation. There have been efforts to build accumulation schemes that overcome the limitations of fixed R1CS. SuperNova [KS22] enables selecting the appropriate R1CS instance at runtime without a recursive circuit that is linear in all R1CS instances. The approach, however, still has limitations. The recursive circuit still requires many (though a constant number of) hashes and a hash-to-group gadget, and additionally, the accumulator, and thus the final proof, is still linear in the total size of all instances. Sangria [Moh23] describes an accumulation scheme for a Plonk-like [GWC19] constraint system with degree-2 gates. It also proposes a solution for higher-degree gates in the future work section but without security proof. These accumulation schemes are built from simple underlying protocols performing a linear combination between an accumulator and a proof. However, the constructions seem ad hoc and need individual security proof. This leads us to our main research questions:

Recipe for accumulation Is there a general recipe for building accumulation schemes?
 Can we formalize this recipe, simplifying the task of constructing secure and efficient accumulation schemes?

Efficient accumulation for ZK-EVM Can we build an accumulation/folding scheme for a language that combines the benefits of the most advanced proof systems today?
 Can we support multiple circuits, high-degree, and lookup gates?

We answer both questions positively. Firstly we show a general compiler that takes any $(2k - 1)$ -move special-sound interactive argument for an NP-complete relation \mathcal{R}_{NP} with an algebraic degree d verifier and construct an efficient IVC-scheme from it. This is done in 4 simple steps.

1. We compress the prover message by committing to them in a homomorphic commitment scheme.
2. Then we apply the Fiat-Shamir transform to yield a secure NARK. [AFK22; Wik21]
3. We build a simple and efficient accumulation scheme that samples a random challenge α and takes a linear combination between the current accumulator and the new NARK.
4. We apply the compiler by [BCLMS21] to yield a secure IVC scheme.

The recursive circuit of this transformation is dominated by only $d + k - 1$ scalar multiplications in the additive group of the commitment scheme² for a protocol with k prover messages and a degree d verifier. For R1CS, where $k = 1$ and $d = 2$, this yields the same protocol and efficiency as Nova[KST22]. We can further reduce the size of the recursive circuit to only $k + 2$ group scalar multiplication, by compressing all verification equations using a random linear combination.

Efficient simple protocols for $\mathcal{R}_{\text{mplkup}}$. Equipped with this compiler, we design PROTOSTAR, a simple and efficient IVC scheme for a highly expressive language $\mathcal{R}_{\text{mplkup}}$ that supports multiple non-uniform circuits and enables high degree and lookup gates. The schemes can be instantiated from any linearly homomorphic vector commitment, e.g., the discrete logarithm-based Pedersen commitment [Ped92], and do not require a trusted setup or the computation of large FFTs. The protocol has several advantages over prior schemes:

Non-uniform IVC without overhead. Each iteration has a program counter \mathbf{pc} that selects one out of I circuits. Part of the circuit constrains \mathbf{pc} ; e.g., \mathbf{pc} could depend on the iteration or indicate which instruction within a VM is executed. The IVC-prover, including the recursive statement, only requires one exponentiation per non-zero bit in the witness. The prover’s computation is independent of I .

Flexible high degree gates. Our protocol supports Plonk-like constraint systems with degree d gates instead of just addition and multiplication. The recursive statement consists of 3 group scalar multiplications and $d + O(1)$ hash and field operations. Unlike in traditional Plonk, there is no additional cost for additional gate types (of degree less than d) and additional selectors. This enables a high level of non-uniformity, even within a circuit.

Lookups, linear and independent of table size. PROTOSTAR supports lookup gates that ensure a value is in some precomputed table T . In each computation step, the prover commits to 2 vectors of length ℓ_{lk} , where ℓ_{lk} is the number of lookups. The prover, in each step, is independent of the table size (assuming free indexing in T). We also support tables that store tuples of size v using 1 additional challenge computations within the recursive circuit.

Our protocols are built of multiple small building blocks. In the protocol for high-degree gates, the prover simply sends the witness, and the degree d verifier checks the circuit with degree d gates. For lookup, we leverage an insight by Haböck [Hab22] on logarithmic derivatives. This yields a protocol where a prover performing ℓ_{lk} in a table of size T only needs to commit to two vectors of length ℓ_{lk} , independent of T . This is the most efficient lookup protocol today. While the verification is linear time, it is low degree (2) and thus compatible

²When instantiated with elliptic curve Pedersen commitments, this translates to $d + k - 1$ elliptic curve multiplications. This is usually the largest component of the recursive statement.

	PROTOSTAR	HyperNova	SuperNova
Language	Degree d Plonk/CCS	Degree d CCS	R1CS (degree 2)
Non-uniform	yes	no	yes
P native	$ \mathbf{w} \mathbb{G}$ $O(\mathbf{w} d \log^2 d) \mathbb{F}$	$ \mathbf{w} \mathbb{G}$ $O(\mathbf{w} d \log^2 d) \mathbb{F}$	$ \mathbf{w} \mathbb{G}$
extra P native w/ lookup	$O(\ell_{\text{lk}}) \mathbb{G}$	$O(T) \mathbb{F}$	N/A
P recursive	$3\mathbb{G}$ $(d + O(1))\mathbb{H} + \mathbb{H}_{\text{in}}$ $(d + O(1)) \mathbb{F}$	$1\mathbb{G}$ $d \log n \mathbb{H} + \mathbb{H}_{\text{in}}$ $O(d \log n) \mathbb{F}$	$2\mathbb{G}$ $\mathbb{H}_{\text{in}} + O_\lambda(1)\mathbb{H} + 1\mathbb{H}_{\mathbb{G}}$
extra P recursive w/ lookup	$1\mathbb{H}$	$O(\log T) \mathbb{H}$ $O(\ell_{\text{lk}} \log T) \mathbb{F}$	N/A

Table 1: The comparison between IVCs.

with our generic compiler. Combining all these yields PROTOSTAR, a new IVC-scheme for $\mathcal{R}_{\text{mplkup}}$. We compare PROTOSTAR, with SuperNova [KS22] and HyperNova [KS23], in Table 1 (for more detail see Corollary 1): In the table, $|\mathbf{w}|$ is the number of non-zero entries of the witness for circuit i , and ℓ_{lk} is the number of lookups in a table of size T . \mathbb{G} is the cost of a group scalar multiplication. \mathbb{F} is the cost of a field multiplication. $d\mathbb{H}$ denotes the cost of hashing d λ -bit numbers. We assume that the cost scales linearly with the size of the input and output. In PROTOSTAR d field elements are hashed once and in HyperNova d field elements are hashed $\log(n)$ times. $\mathbb{H}_{\mathbb{G}}$ is the cost of a hash-to-group function. \mathbb{H}_{in} is the cost of hashing the public input and the accumulator instance. Note that the $O_\lambda(1)\mathbb{H}$ in SuperNova’s recursive circuit involves constant number of hashes to the input of two accumulator instances and one circuit verification key, by using multiset-based offline memory checking in a circuit [SAGL18].

Concurrent work. In a paper concurrent with this work, Kothapalli and Setty [KS23] introduce an IVC for high degree relations. They use a generalization of R1CS called customizable constraint systems (CCS) [STW23] that covers the Plonkish relations. It also enables gates with a high additive fan-in. PROTOSTAR also has no restriction to the fan-in an individual gate has, but we subsequently showed that our compiler can also be directly applied to CCS (Appendix B). HyperNova is based on so-called multi-folding schemes. They also provide a lookup argument suitable for recursive arguments. However, they do not explicitly explain how to integrate lookup to Plonk/CCS in their IVC scheme or provide any explicit constructions for non-uniform computations. Their scheme is built using sumchecks [LFKN92] and the resulting IVC recursive circuit is dominated by 1 group scalar multiplication, $d \log n + \ell_{\text{in}}$ hash operations and $O(d \log n + \ell_{\text{in}})$ field multiplications where d is the custom gate degree, n is the number of gates and ℓ_{in} is the public input length.

In comparison, our IVC recursive circuit, even with lookup and non-uniformity support, is only dominated by 3 group scalar multiplications and $O(\ell_{\text{in}} + d)$ field/hash operations, entirely independent of n . The 2 additional group operations compared to HyperNova are likely offset by the additional lookup support [XCZBFKC22] and the significantly fewer hashes and non-native field operations (d vs. $d \log(n)$). A detailed comparison is given in Table 1.

For a lookup relation with table size T and ℓ_{lk} lookup gates, their accumulation/folding scheme leads to an accumulation prover whose work is dominated by $O(T)$ field operations and an accumulation verifier whose work is dominated by $O(\ell_{\text{lk}} \log T)$ field operations and $O(\log T)$ hashes. This is undesirable when the table size $T \gg \ell_{\text{lk}}$. In comparison, our scheme has prover complexity $O(\ell_{\text{lk}})$ and the verifier is only dominated by 3 group scalar multiplications, 2 hashes and 2 field multiplications. Moreover, the lookup support adds almost no overhead to the IVC scheme for high-degree Plonk relations. In particular, it adds no group scalar multiplications. Lastly, their lookup scheme does not support vector-valued lookups, which is essential for applications like ZK-EVM and encoding bit-wise operations in circuits.

1.1 Technical overview

Given an NP-complete relation \mathcal{R} , we introduce a generic framework for constructing efficient incremental verifiable computation (IVC) schemes with predicates expressed in \mathcal{R} . For \mathcal{R} being the non-uniform Plonkup circuit satisfiability relation, we obtain an efficient (non-uniform) IVC scheme for proving correct program executions on stateful machines (e.g., EVM). The framework starts by designing a simple special-sound protocol Π_{sps} for relation \mathcal{R} , which is easy to analyze. Next, we use a generic compiler to transform Π_{sps} into a Non-interactive Argument of Knowledge Scheme (NARK) whose verification predicate is easy to accumulate/fold. Finally, we build an efficient accumulation/folding scheme for the NARK verifier, and apply the generic compiler from [BCLMS21] to obtain the IVC/PCD scheme for relation \mathcal{R} . We describe the workflow in Figure 1.

The paper begins by describing the compiler from special-sound protocols to NARKs in Section 3, and presents an efficient accumulation scheme for the compiled NARK verifier in Section 3.4. Next, we describe simple and efficient special-sound protocols for Plonkup circuit-satisfiability relations in Section 5 and extend it to support non-uniform computation in Section 6. Similarly, we extend the CCS relation [STW23] to support non-uniform computation and lookup in Appendix B. We give an overview of our approach below.

Efficient IVCs from special-sound protocols. Let Π_{sps} be any *multi-round* special-sound protocol for some relation \mathcal{R} , in which the verifier is *algebraic*, that is, the verifier algorithm only checks algebraic equations over the input and the prover messages. E.g., the following naive protocol for the Hadamard product relation over vectors $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{F}^n$ is special-sound and has a degree-2 algebraic verifier: The prover simply sends the vectors

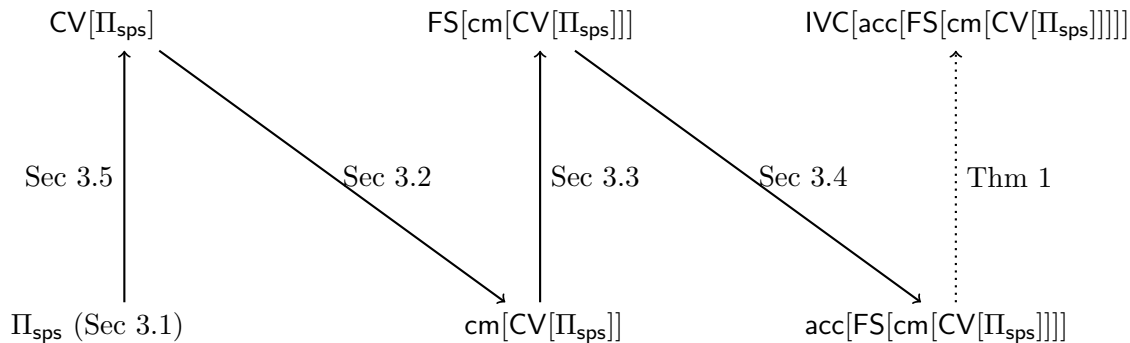


Figure 1: The workflow for building an IVC from a special sound protocol. We start from a special-sound protocol Π_{sps} for an NP-complete relation \mathcal{R}_{NP} , and transform it to $\text{CV}[\Pi_{\text{sps}}]$ with a compressed verifier check. $\text{CV}[\Pi_{\text{sps}}]$ is converted to a NARK $\text{FS}[\text{cm}[\text{CV}[\Pi_{\text{sps}}]]]$ via commit-and-open and the Fiat-Shamir transform. We then build a generic accumulation scheme for the NARK and apply Theorem 1 from [BCLMS21] to obtain the IVC scheme. This last connection is dotted as it requires heuristically replacing random oracles with cryptographic hash functions.

\mathbf{a} , \mathbf{b} , \mathbf{c} to the verifier, and the verifier checks that $a_i \cdot b_i = c_i$ for all $i \in [n]$. However, as shown in the example, the prover message can be large in Π_{sps} and the folding scheme can be expensive if we directly accumulate the verifier predicate. Inspired by the splitting accumulation scheme [BCLMS21], to enable efficient accumulation/folding, we split each prover message into a short instance and a large opening, where the short instance is built from the homomorphic commitment to the prover message. Next, we use the Fiat-Shamir transform to compile the protocol into a NARK where the verifier challenges are generated from a random oracle.

Now we can view the NARK transcript as an accumulator (or a relaxed NP instance-witness pair in the language of folding schemes), where the accumulator instance consists of the prover message commitments and the verifier challenges; while the accumulator witness consists of the prover messages (i.e., the opening to the commitments). Note we also need to introduce an error vector/commitment into the accumulator witness/instance to absorb the “noise” that arises after each accumulation/folding step.

In the accumulation scheme, given two accumulators (or NARK proofs), the prover folds the witnesses and the instances of both accumulators via a random linear combination and generates a list of d “error-correcting terms” as accumulation proof (d is the degree of the NARK verifier); the verifier only needs to check that the folded accumulator instance is consistent with the accumulation proof and the original instances being folded, both of which are small. After finishing all the accumulation steps, a decider applies a final check to the accumulator, scrutinizing that (i) the accumulator witness is consistent with the commitments in the accumulator instance, and (ii) the “relaxed” NARK verifier check still

passes. Here by “relaxed” we mean that the algebraic equation also involves the error vector in the accumulator. If the decider accepts, this implies that all accumulated NARKs were valid and thus that all accumulated statements are in \mathcal{R} (and the prover knows witnesses for these statements).

Finally, given the accumulation scheme, if the relation \mathcal{R} is NP-complete, we can apply the compiler in [BCLMS21] to obtain an efficient IVC scheme with predicates expressed in \mathcal{R} .

In Theorem 3, we show that for any $(2k-1)$ -move³ special-sound protocols with degree- d verifiers, the resulting IVC recursive circuit only involves $k + d + O(1)$ hashes, $k + 1$ non-native field operations and $k + d - 1$ commitment group scalar multiplications. We also introduce a generic approach for further reducing the number of group operations to $k + 2$ in Section 3.5. This is favorable for $d \geq 3$. The idea is to compress all ℓ degree d verification checks into a single verification check using a random linear combination with powers of a challenge β . This means that error-correcting terms are field elements and, thus, can be sent directly without committing to them. The prover also sends a single commitment to powers of β and powers of $\beta^{\sqrt{\ell}}$. The verification equation uses one power of β and one power of $\beta^{\sqrt{\ell}}$, which increases the degree of the verification check to $d + 2$. The verifier also checks the correctness of the powers of β using $2\sqrt{\ell}$ degree 2 checks.

Special-sound protocols for (non-uniform) Plonkup relations. Given the generic compiler above, our ultimate goal of constructing a (non-uniform) IVC scheme for zkEVM becomes much easier. It is now sufficient to design a multi-round special-sound protocol for the (non-uniform) Plonkup relation. We describe the components of the special-sound protocol in Figure 2. Note we also extend CCS relation [STW23] to support lookup and non-uniform computation and build a special-sound protocol for it (See Figure 2). Recall that a Plonkup circuit-satisfiability relation consists of three modular relations, namely, (i) a high-degree gate relation checking that each custom gate is satisfied; (ii) a permutation (wiring-identity) relation checking that different gate values are consistent if the same wire connects them, and (iii) a lookup relation checking that a subset of gate values belongs to a preprocessed table. The special-sound protocols for the permutation and high-degree gate relations are trivial, where the prover directly sends the witness to the verifier, and the verifier checks that the permutation/high-degree gate relation holds. The degree of the permutation check is only 1, and the degree of the gate-check is the highest degree in the custom gate formula.

The special-sound protocol for the lookup relation \mathcal{R}_{LK} is more interesting as the statement of the lookup relation is not algebraic. Inspired by the log-derivative lookup scheme [Hab22], in Section 4.3, we design a simple 3-move special-sound protocol Π_{LK} for \mathcal{R}_{LK} , in which the verifier degree is only 2. A great feature of Π_{LK} is that the number of non-zero elements in the prover messages is only proportional to the number of lookups,

³ k prover messages, $k - 1$ challenges

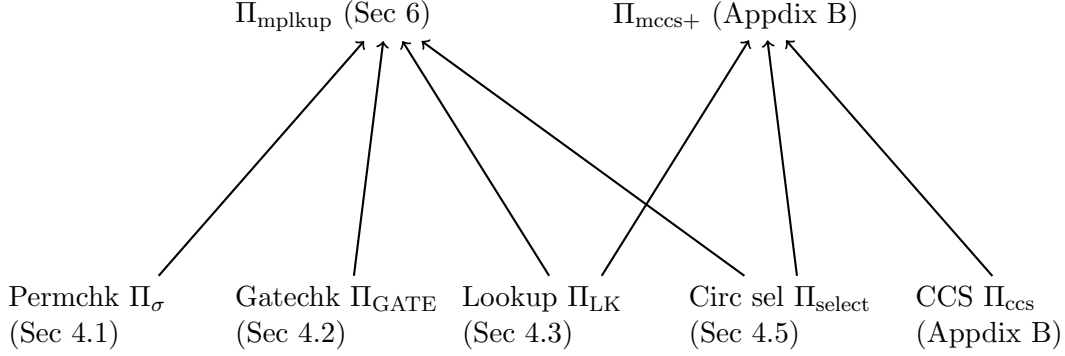


Figure 2: The special-sound protocols for PROTOSTAR and $\text{PROTOSTAR}_{\text{CCS}}$. The special-sound protocol Π_{mplkup} for the multi-circuit Plonkup relation $\mathcal{R}_{\text{mplkup}}$ consists of the sub-protocols for permutation, high-degree custom gate, lookup, and circuit selection relations. The special-sound protocol $\Pi_{\text{mccs+}}$ for the extended CCS relation $\mathcal{R}_{\text{mccs+}}$ consists of the sub-protocols for lookup, circuit selection, as well as the CCS relation [STW23]. From Π_{mplkup} or $\Pi_{\text{mccs+}}$, we can apply the workflow described in Fig 1 to obtain the IVC schemes PROTOSTAR or $\text{PROTOSTAR}_{\text{CCS}}$.

but independent of the table size. Thus the IVC prover complexity for computing the prover message commitments is independent of the table size, which is advantageous when the table size is much larger than the witness size. However, the prover work for computing the error terms is not independent of the table size because the accumulator is not sparse. Fortunately, we observe that the prover can efficiently update the error term commitments without recomputing the error term vectors from scratch, thus preserving the efficiency of the accumulation prover. Moreover, we extend Π_{LK} in Section 4.4 to further support vector-valued lookup, where each table entry is a vector of elements. This feature is useful in applications like zkEVM and for simulating bit operations in circuits.

Given the special-sound protocols for permutation/high-degree gate/lookup relations, the special-sound protocol Π_{plonkup} for Plonkup is just a parallel composition of the three protocols. Furthermore, in Section 6, we apply a simple trick to support *non-uniform* IVC. More precisely, let $\{\mathcal{C}_i\}_{i=1}^I$ be I different branch circuits (e.g., the set of supported instructions in EVM), let $\mathbf{pi} := (pc, \mathbf{pi}')$ be the public input where $pc \in [I]$ is a program counter indicating which instruction/branch circuit is going to be executed in the next IVC step. Our goal is to prove that $(\mathbf{pi}, \mathbf{w})$ is in the relation $\mathcal{R}_{\text{mplkup}}$ in the sense that $\mathcal{C}_{pc}(\mathbf{pi}, \mathbf{w}) = 0$ for witness \mathbf{w} . The relation statement can also add additional constraints on pc depending on the applications. The special-sound protocol for $\mathcal{R}_{\text{mplkup}}$ is almost identical to Π_{plonkup} for the Plonkup relation, except that the prover further sends a bool vector $\mathbf{b} \in \mathbb{F}^I$, and the verifier uses $2I$ degree 2 equations to check that $b_{pc} = 1$ and $b_i = 0 \forall i \neq pc$. Additionally, each algebraic equation \mathcal{G} checked in Π_{plonkup} is replaced with

$\sum_{i=1}^I \mathcal{G}_i \cdot b_i$ where \mathcal{G}_i ($1 \leq i \leq I$) is the corresponding gate in the i -th branch circuit. The resulting special-sound protocol has 3 moves, and the verifier degree is $d + 1$, where d is the highest degree of the custom gates. This means that the IVC scheme for the non-uniform Plonkup relation adds negligible overhead to that for the Plonkup relation.

2 Preliminaries

2.1 Special-sound Protocols and Fiat-Shamir Transform

We define special-soundness and non-interactive arguments according to the definitions by [AFK22].

Definition 1 (Public-coin interactive proof). *An interactive proof $\Pi = (\mathsf{P}, \mathsf{V})$ for relation \mathcal{R} is an interactive protocol between two probabilistic machines, a prover P , and a polynomial time verifier V . Both P and V take as public input a statement \mathbf{pi} and, additionally, P takes as private input a witness $\mathbf{w} \in \mathcal{R}(\mathbf{pi})$. The verifier V outputs 0 if it accepts and a non-zero value otherwise. Its output is denoted by $(\mathsf{P}(\mathbf{w}), \mathsf{V})(\mathbf{pi})$. Accordingly, we say the corresponding transcript (i.e., the set of all messages exchanged in the protocol execution) is accepting or rejecting. The protocol is public coin if the verifier randomness is public. The verifier messages are referred to as challenges. Π is a $(2k - 1)$ -move protocol if there are k prover messages and $k - 1$ verifier messages.*

Definition 2 (Tree of transcript). *Let $\mu \in \mathbb{N}$ and $(a_1, \dots, a_\mu) \in \mathbb{N}^\mu$. An (a_1, \dots, a_μ) -tree of transcript for a $(2\mu + 1)$ -move public-coin interactive proof Π is a set of $a_1 \cdot a_2 \cdots a_\mu$ accepting transcripts arranged in a tree of depth μ and arity a_1, \dots, a_μ respectively. The nodes in the tree correspond to the prover messages and the edges to the verifier's challenges. Every internal node at depth $i - 1$ ($1 \leq i \leq \mu$) has a_i children with distinct challenges. Every transcript corresponds to one path from the root to a leaf node. We simply write the transcripts as an (a^μ) -tree of transcript when $a = a_1 = a_2 = \dots = a_\mu$.*

Definition 3 (Special-sound Interactive Protocol). *Let $\mu, N \in \mathbb{N}$ and $(a_1, \dots, a_\mu) \in \mathbb{N}^\mu$. A $(2\mu + 1)$ -move public-coin interactive proof Π for relation \mathcal{R} where the verifier samples its challenges from a set of size N is (a_1, \dots, a_μ) -out-of- N special-sound if there exists a polynomial time algorithm that, on input \mathbf{pi} and any (a_1, \dots, a_μ) -tree of transcript for Π outputs $\mathbf{w} \in \mathcal{R}(\mathbf{pi})$. We simply denote the protocol as an a^μ -out-of- N (or a^μ) special-sound protocol if $a = a_1 = a_2 = \dots = a_\mu$.*

Definition 4 (Random-Oracle Non-Interactive Argument of Knowledge (RO-NARK)). *A non-interactive random oracle proof for relation \mathcal{R} is a pair (P, V) of probabilistic random-oracle algorithms, such that: Given $(\mathbf{pi}, \mathbf{w}) \in \mathcal{R}$ and access to a random oracle ρ_{NARK} , the prover $\mathsf{P}^{\rho_{\text{NARK}}}(\mathbf{pi}, \mathbf{w})$ outputs a proof π . Given \mathbf{pi} , a proof π , and access to the same random oracle ρ_{NARK} , the verifier $\mathsf{V}^{\rho_{\text{NARK}}}(\mathbf{pi}, \pi)$ outputs 0 to accept or any other value to reject.*

Perfect Completeness: *The NARK has perfect completeness if for all $(\mathbf{pi}, \mathbf{w}) \in \mathcal{R}$*

$$P[V^{\rho_{\text{NARK}}}(\mathbf{pi}, P^{\rho_{\text{NARK}}}(\mathbf{pi}, \mathbf{w})) = 0] = 1$$

Knowledge Soundness: *The NARK has adaptive knowledge-soundness with knowledge error $\kappa : \mathbb{N} \times \mathbb{N} \rightarrow [0, 1]$ if there exists a knowledge extractor Ext , with the following properties: The extractor, given input n , and oracle-access to any polynomial-time Q -query random oracle prover P^* that outputs statement of size n , runs in an expected polynomial time in $|\mathbf{pi}| + Q$, and outputs $\{(\mathbf{pi}, \pi, \mathbf{aux}, v; \mathbf{w})\}$ such that a) $(\mathbf{pi}, \pi, \mathbf{aux}, v)$ is identically distributed to $\{(\mathbf{pi}, \pi, \mathbf{aux}, v) : (\mathbf{pi}, \pi, \mathbf{aux}) \leftarrow P^{*, \rho_{\text{NARK}}}, v \leftarrow V^{\rho_{\text{NARK}}}(\mathbf{pi}, \pi)$ and b)*

$$\Pr \left[\begin{array}{l} (\mathbf{pi}; \mathbf{w}) \in \mathcal{R} \\ V^{\rho_{\text{NARK}}}(\mathbf{pi}, \pi) = 0 \end{array} : \{(\mathbf{pi}, \pi, \mathbf{aux}, v; \mathbf{w})\} \leftarrow \text{Ext}^{P^*} \right] \geq \frac{\epsilon(P^*) - \kappa(n, Q)}{\text{poly}(n)},$$

where $\epsilon(P^*)$ is P 's success probability, i.e. $\epsilon(P^*) = P[V^{\rho_{\text{NARK}}}(\mathbf{pi}, \pi) = 0 : (\mathbf{pi}, \pi) \leftarrow P^{*, \rho_{\text{NARK}}}]$. Here, Ext implements ρ_{NARK} for P^* ; in particular, it can arbitrarily program the random oracle.

Definition 5 (Fiat-Shamir Transform (adaptive)). *The Fiat-Shamir transform $\text{FS}[\Pi] = (P_{\text{fs}}, V_{\text{fs}})$ is a RO-NARK, where $P^{\rho_{\text{NARK}}}(\mathbf{pi}; \mathbf{w})$ runs $P(\mathbf{pi}; \mathbf{w})$ but instead of receiving challenge c_i , on message m_i , from the verifier, it computes it as follows:*

$$c_i = \rho_{\text{NARK}}(c_{i-1}, m_i) \tag{1}$$

and $c_0 = \rho_{\text{NARK}}(\mathbf{pi})$. $P_{\text{fs}}^{\rho_{\text{NARK}}}$ outputs $\pi = (m_1, \dots, m_\mu)$. The verifier $V_{\text{fs}}^{\rho_{\text{NARK}}}$ accepts, if V accepts the transcript $(m_1, c_1, \dots, m_\mu, c_\mu, m_{\mu+1})$ for input \mathbf{pi} and the challenges are computed as per equation (1).

2.2 Adaptive Fiat-Shamir transform

Lemma 1 (Fiat-Shamir transform of Special-sound Protocols [AFK22]). *The Fiat-Shamir transform of a $(\alpha_1, \dots, \alpha_\mu)$ -out-of- N special-sound interactive proof Π is knowledge sound with knowledge error*

$$\kappa_{\text{fs}}(Q) = (Q + 1)\kappa$$

where $\kappa = 1 - \prod(1 - \frac{\alpha_i}{N})$ is the knowledge error of the interactive proof Π .

2.3 Commitment Scheme

Definition 6 (Commitment Scheme). $\text{cm} = (\text{Setup}, \text{Commit})$ is a binding commitment scheme, consisting of two algorithms:

$\text{Setup}(1^\lambda) \rightarrow \text{ck}$ takes as input the security parameter and outputs a commitment key ck .

$\text{Commit}(\text{ck}, \mathbf{m} \in \mathcal{M}) \rightarrow C \in \mathcal{C}$, takes as input the commitment key ck and a message \mathbf{m} in

\mathcal{M} and outputs a commitment $C \in \mathcal{C}$.

The scheme is binding if for all polynomial-time randomized algorithms P^* :

$$\Pr \left[\begin{array}{c} \text{Commit}(\text{ck}, \mathbf{m}) = \text{Commit}(\text{ck}, \mathbf{m}') \\ \wedge \\ \mathbf{m} \neq \mathbf{m}' \end{array} \middle| \begin{array}{c} \text{ck} \leftarrow \text{Setup}(1^\lambda) \\ \mathbf{m}, \mathbf{m}' \leftarrow P^*(\text{ck}) \end{array} \right] = \text{negl}(\lambda)$$

Homomorphic commitment. We say the commitment is homomorphic if $(\mathcal{C}, +)$ is an additive group of prime order p .

2.4 Incremental Verifiable Computation (IVC)

We adapt and simplify the definition from [BCLMS21; KST22].

Definition 7 (IVC). An incremental verifiable computation (IVC) scheme for function predicates expressed in a circuit-satisfiability relation \mathcal{R}_{NP} is a tuple of algorithms $\text{IVC} = (P_{\text{IVC}}, V_{\text{IVC}})$ with the following syntax and properties:

- $P_{\text{IVC}}(m, z_0, z_m, z_{m-1}, \mathbf{w}_{\text{loc}}, \pi_{m-1}) \rightarrow \pi_m$. The IVC prover P_{IVC} takes as input a program output z_m at step m , local data \mathbf{w}_{loc} , initial input z_0 , previous program output z_{m-1} and proof π_{m-1} and outputs a new IVC proof π_m .
- $V_{\text{IVC}}(m, z_0, z_m, \pi_m) \rightarrow b$. The IVC verifier V_{IVC} takes the initial input z_0 , the output z_m at step m , and an IVC proof π_m , ‘accepts’ by outputting $b = 0$ and ‘rejects’ otherwise.

The scheme IVC has perfect adversarial completeness if for any function predicate ϕ expressible in \mathcal{R}_{NP} , and any, possibly adversarially created, $(m, z_0, z_m, z_{m-1}, \mathbf{w}_{\text{loc}}, \pi_{m-1})$ such that

$$\phi(z_0, z_m, z_{m-1}, \mathbf{w}_{\text{loc}}) \wedge (V_{\text{IVC}}(m-1, z_0, z_{m-1}, \pi_{m-1}) = 0)$$

it holds that $V_{\text{IVC}}(m, z_0, z_m, \pi_m)$ accepts for proof $\pi_m \leftarrow P_{\text{IVC}}(m, z_0, z_{m-1}, z_m, \mathbf{w}_{\text{loc}}, \pi_{m-1})$.

The scheme IVC has knowledge soundness if for every expected polynomial-time adversary P^* , there exists an expected polynomial-time extractor Ext_{P^*} such that

$$\Pr \left[\begin{array}{c} V_{\text{IVC}}(m, z_0, z, \pi_m) = 0 \wedge \\ (\exists i \in [m], \neg \phi(z_0, z_i, z_{i-1}, \mathbf{w}_i)) \\ \vee z \neq z_m \end{array} \middle| \begin{array}{c} [\phi, (m, z_0, z, \pi_m)] \leftarrow P^* \\ [z_i, \mathbf{w}_i]_{i=1}^m \leftarrow \text{Ext}_{P^*} \end{array} \right] \leq \text{negl}(\lambda).$$

Here m is a constant.

2.5 Simple Accumulation

We take definitions and proofs from [BCLMS21].

Definition 8 (Accumulation Scheme). *An accumulation scheme for a NARK (P_{NARK}, V_{NARK}) is a triple of algorithms $\text{acc} = (P_{\text{acc}}, V_{\text{acc}}, D)$, all of which have access to the same random oracle ρ_{acc} as well as ρ_{NARK} , the oracle for the NARK. The algorithms have the following syntax and properties:*

- $P_{\text{acc}}(\text{pi}, \pi = (\pi.x, \pi.w), \text{acc} = (\text{acc}.x, \text{acc}.w)) \rightarrow \{\text{acc}' = (\text{acc}'.x, \text{acc}'.w), \text{pf}\}$. *The accumulation prover P_{acc} takes as input a statement pi , NARK proof π , and an accumulator acc and outputs a new accumulator acc' and correction terms pf .*
- $V_{\text{acc}}(\text{pi}, \pi.x, \text{acc}.x, \text{acc}'.x, \text{pf}) \rightarrow v$. *The accumulation verifier takes as input the statement pi , the instances of the NARK proof, the old and new accumulator, the correction terms, and ‘accepts’ by outputting 0 and ‘rejects’ otherwise.*
- $D(\text{acc}) \rightarrow v$. *The decider on input acc ‘accepts’ by outputting 0 and ‘rejects’ otherwise.*

An accumulation scheme has knowledge-soundness with knowledge error κ if the RO-NARK (P', V') has knowledge error κ for the relation

$$\mathcal{R}_{\text{acc}}((\text{pi}, \pi.x, \text{acc}.x); (\pi.w, \text{acc}.w)) : (V_{NARK}(\text{pi}, \pi) = 0 \wedge D(\text{acc}) = 0) ,$$

where P' outputs acc' , pf and V' on input $((\text{pi}, \pi.x, \text{acc}.x), (\text{acc}', \text{pf}))$ accepts if $D(\text{acc}')$ and $V_{\text{acc}}(\text{pi}, \pi.x, \text{acc}.x, \text{acc}'.x, \text{pf}) = 0$.

The scheme has perfect completeness if the RO-NARK (P', V') has perfect completeness for \mathcal{R}_{acc} .

Theorem 1 (IVC from accumulation[BCLMS21]). *Given a standard-model NARK for circuit-satisfiability and a standard-model accumulation scheme (Definition 8) for that NARK, both with negligible knowledge error, there exists an efficient transformation that outputs an IVC scheme (see Section 3.2 of [BCLMS21]) for constant-depth compliance predicates, assuming that the circuit complexity of the accumulation verifier V_{acc} is sub-linear in its input.*

Random Oracle. Note that both the NARK and accumulation scheme we construct are in the random oracle model. However, Theorem 1 requires a NARK and an accumulation scheme in the standard model. It remains an open problem to construct such schemes. However, we can heuristically instantiate the random oracle with a cryptographic hash function and assume that the resulting schemes still have knowledge soundness.

Definition 9 (Fiat-Shamir Heuristic). *The Fiat-Shamir Heuristic, relative to a secure cryptographic hash function H , states that a random oracle NARK with negligible knowledge error yields a NARK that has negligible knowledge error in the standard (CRS) model if the random oracle is replaced with H .*

Complexity. The IVC transformation from [BCLMS21] recursively proves that the accumulation was performed correctly. To do that, it implements V_{acc} as a circuit and proves that the previous accumulation step was done correctly. Note that this recursive circuit is independent of the size of $\pi, \mathbf{w}, \text{acc}, \mathbf{w}$ and the runtime of D . The IVC prover is linear in the size of the recursive circuit plus the size of the IVC computation step expressed as a circuit. The final IVC verifier and the IVC proof size are linear in these components. This can be reduced using an additional SNARK as in [KST22].

PCD. IVC can be generalized to arbitrary DAGs instead of just path graphs in a primitive called proof-carrying data [BCCT13]. Accumulation schemes can be compiled into full PCD if they support accumulating an arbitrary number of accumulators and proofs [BCMS20; BCLMS21]. For simplicity, we only build accumulation for one proof and one accumulator, as well as for two accumulators. This enables PCD for DAGs of degree two. By transforming higher degree graphs into degree two graphs (by converting each degree d node into a $\log_2(d)$ depth tree), we can achieve PCD for these graphs.

Outsourcing the decider. In the accumulation to IVC transformation, the IVC proof is linear in the accumulator, and the IVC verifier runs the decider. The accumulation schemes we construct are linear in the witness of a single computation step. However, we can outsource the decider by providing a SNARK that, given $\text{acc}.x$, proves knowledge of $\text{acc}.\mathbf{w}$, such that $D(\text{acc}) = 0$. Nova [KST22] constructs a custom, concretely efficient SNARK for their accumulation/folding scheme. However, when outsourcing the decider, the IVC cannot continue. This breaks the strict completeness requirement of IVC, which says that any prover can continue from any valid IVC proof. Nevertheless, this may be fine for some applications of IVC.

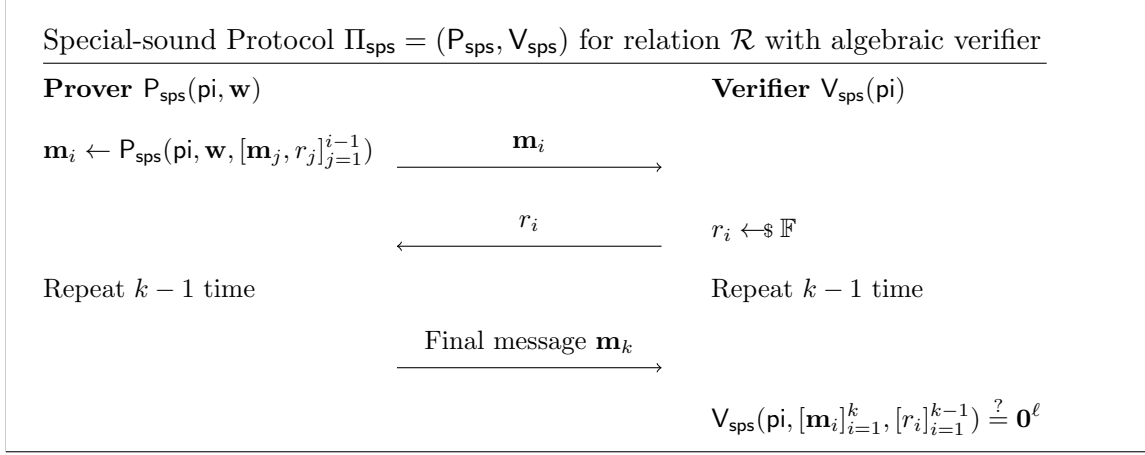
3 Protocols

3.1 Special-sound Protocols

In this section, we describe a class of special-sound protocols whose verifier is algebraic. The protocol Π_{sps} has 3 essential parameters $k, d, \ell \in \mathbb{N}$, meaning that Π_{sps} is a $(2k - 1)$ -move protocol with verifier degree d and output length ℓ (i.e. the verifier checks ℓ degree d algebraic equations). In each round i ($1 \leq i \leq k$), the prover $P_{\text{sps}}(\mathbf{pi}, \mathbf{w}, [\mathbf{m}_j, r_j]_{j=1}^{i-1})$ generates the next message \mathbf{m}_i on input the public input \mathbf{pi} , the witness \mathbf{w} , and the current transcript $[\mathbf{m}_j, r_j]_{j=1}^{i-1}$, and sends \mathbf{m}_i to the verifier; the verifier replies with a random challenge $r_i \in \mathbb{F}$. After the final message \mathbf{m}_k , the verifier computes the algebraic map V_{sps} and checks that the output is a zero vector of length ℓ . More precisely, $\deg(V_{\text{sps}}) = d$, s.t.

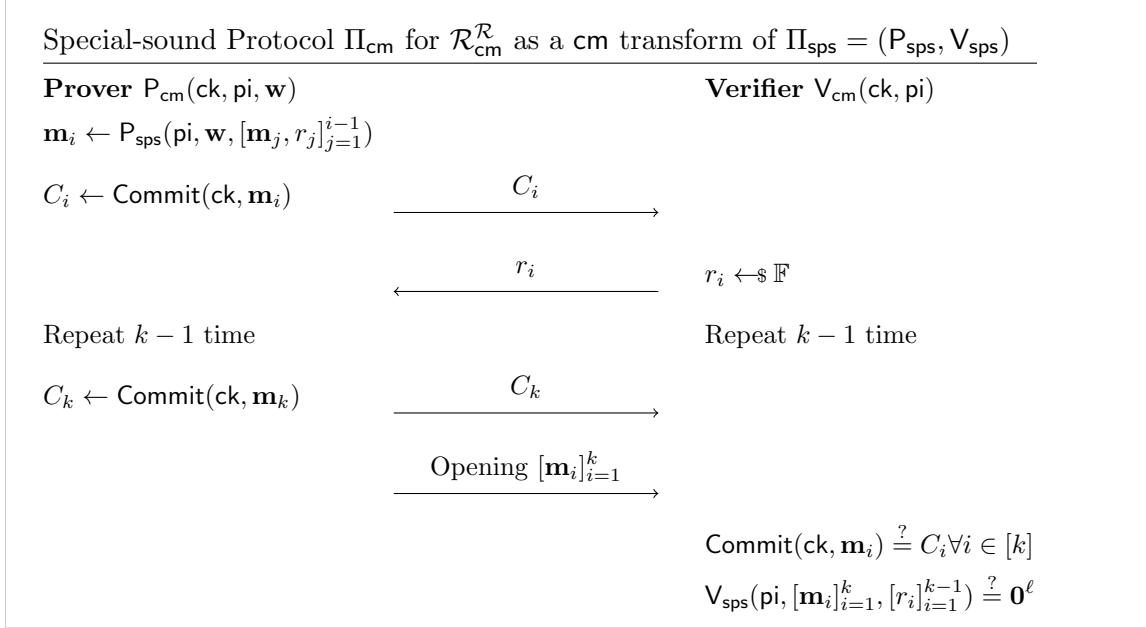
$$V_{\text{sps}}(\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}) := \sum_{j=0}^d f_j^{V_{\text{sps}}}(\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}),$$

where $f_j^{\mathcal{V}_{\text{sps}}}$ is a homogeneous degree- j algebraic map that outputs a vector of ℓ field elements. We describe the special-sound protocol Π_{sps} below.



3.2 Commit and Open

For a commitment scheme $\text{cm} = (\text{Setup}, \text{Commit})$, consider the following relation $\mathcal{R}_{\text{cm}}^{\mathcal{R}} = (x; \mathbf{w}, \mathbf{m} \in \mathcal{M}, \mathbf{m}' \in \mathcal{M}) : \{(x, \mathbf{w}) \in \mathcal{R} \vee (\text{Commit}(\mathbf{m}) = \text{Commit}(\mathbf{m}') \wedge \mathbf{m} \neq \mathbf{m}')\}$. The relation's witness is either a valid witness for \mathcal{R} or a break of the commitment scheme cm . We now design a special-sound protocol $\Pi_{\text{cm}} = (\mathcal{P}_{\text{cm}}, \mathcal{V}_{\text{cm}})$ for $\mathcal{R}_{\text{cm}}^{\mathcal{R}}$ given $\Pi_{\text{sps}} = (\mathcal{P}_{\text{sps}}, \mathcal{V}_{\text{sps}})$, a special-sound protocol for \mathcal{R} . \mathcal{P}_{cm} runs \mathcal{P}_{sps} to generate the i th message and then commits to the message. Along with the final message, \mathcal{P}_{cm} sends the opening to the commitment. The verifier \mathcal{V}_{cm} checks the correctness of the commitments and runs \mathcal{V}_{sps} on the commitment openings.



Lemma 2 (Π_{cm} is (a_1, \dots, a_μ) -special-sound). *Let Π_{sps} be an (a_1, \dots, a_μ) -out-of- N special-sound protocol for relation \mathcal{R} , where the prover messages are all in a set \mathcal{M} . Let $(\text{Setup}, \text{Commit})$ be a binding commitment scheme for messages in \mathcal{M} . For $\text{ck} \leftarrow \text{Setup}_{\text{cm}}(1^\lambda)$ let $\mathcal{R}_{\text{cm}} = (\text{pi}; \mathbf{w}, m \in \mathcal{M}, m' \in \mathcal{M}) : (\text{pi}; \mathbf{w}) \in \mathcal{R} \vee (\text{Commit}(\text{ck}, m) = \text{Commit}(\text{ck}, m') \wedge m \neq m')$. Then $\Pi_{\text{cm}} = \text{cm}[\Pi_{\text{sps}}]$ is an (a_1, \dots, a_μ) -out-of- N special-sound protocol for $\mathcal{R}_{\text{cm}}^{\mathcal{R}}$.*

Proof. Let Ext_{sps} be the extractor for Π_{sps} . We will construct Ext_{cm} for Π_{cm} that computes a witness for \mathcal{R}_{cm} , i.e., a witness for \mathcal{R} or a collision for cm given an (a_1, \dots, a_μ) -transcript tree for Π_{cm} . The extractor Ext_{cm} first checks whether there exist two transcripts that have inconsistent final messages. That is, the final message opening is different for the nodes in the intersection of the root-to-leaf paths of these two transcripts. This means we have \mathbf{m}_i and \mathbf{m}'_i , such that $\text{Commit}(\mathbf{m}_i) = \text{Commit}(\mathbf{m}'_i)$ and $\mathbf{m}_i \neq \mathbf{m}'_i$. This is a break for cm , i.e., a valid witness for \mathcal{R}_{cm} . Otherwise Ext_{cm} builds a transcript tree for Π_{sps} by replacing all commitments with the openings and use Ext_{sps} to compute $\mathbf{w} \in \mathcal{R}(\text{pi})$, such that $(\mathbf{w}, \perp, \perp) \in \mathcal{R}_{\text{cm}}(\text{pi})$. \square

3.3 Fiat-Shamir transform

Let ρ_{NARK} be a random oracle. Let Π_{cm} be the commit-and-open protocol for the special-sound protocol $\Pi_{\text{sps}} = (\text{P}_{\text{sps}}, \text{V}_{\text{sps}})$. The Fiat-Shamir Transform $\text{FS}[\Pi_{\text{cm}}]$ of the protocol Π_{cm} is the following. By Lemma 1, $\text{FS}[\Pi_{\text{cm}}]$ is knowledge sound if Π_{sps} is special-sound.

Fiat-Shamir Transform FS of Special-sound Protocol II for relation $\mathcal{R}_{\text{cm}}^{\mathcal{R}}: \text{FS}[\Pi_{\text{cm}}]$	
Prover $P_{\text{NARK}}^{\rho_{\text{NARK}}}(\text{ck}, \text{pi}, \mathbf{w})$ $r_0 \leftarrow \rho_{\text{NARK}}(\text{pi})$ For $i \in [k-1]$: $\mathbf{m}_i \leftarrow P_{\text{sps}}(\text{pi}, \mathbf{w}, [\mathbf{m}_j, r_j]_{j=1}^{i-1})$ $C_i \leftarrow \text{Commit}(\text{ck}, \mathbf{m}_i)$ $r_i \leftarrow \rho_{\text{NARK}}(r_{i-1}, C_i)$ $\mathbf{m}_k \leftarrow P_{\text{sps}}(\text{pi}, \mathbf{w}, [\mathbf{m}_j, r_j]_{j=1}^{k-1})$ $C_k \leftarrow \text{Commit}(\text{ck}, \mathbf{m}_k)$	Verifier $V_{\text{NARK}}^{\rho_{\text{NARK}}}(\text{ck}, \text{pi})$ $r_0 \leftarrow \rho_{\text{NARK}}(x)$ $r_i \leftarrow \rho_{\text{NARK}}(r_{i-1}, C_i) \forall i \in [k-1]$ $\text{Commit}(\text{ck}, \mathbf{m}_i) \stackrel{?}{=} C_i \forall i \in [k]$ $V_{\text{sps}}(\text{pi}, \pi.x, \pi.\mathbf{w}, [r_i]_{i=1}^{k-1}) \stackrel{?}{=} \mathbf{0}^\ell$
$\xrightarrow{\pi.x = [C_i]_{i=1}^k}$ $\xrightarrow{\pi.\mathbf{w} = [\mathbf{m}_i]_{i=1}^k}$	

3.4 Accumulation Scheme for V_{NARK}

Let ρ_{acc} and ρ_{NARK} be two random oracles, and let V_{NARK} be the verifier of $\text{FS}[\Pi_{\text{cm}}]$ in Section 3.3, whose underlying special-sound protocol is $\Pi_{\text{sps}} = (P_{\text{sps}}, V_{\text{sps}})$ for a relation \mathcal{R} . We describe the accumulation scheme for V_{NARK} .

The accumulated predicate. The predicate to be accumulated is the “relaxed” verifier check of the NARK scheme $\text{FS}[\Pi_{\text{cm}}]$ for relation \mathcal{R} . Namely, given public input $\text{pi} \in \mathcal{M}^{\ell_{\text{in}}}$, random challenges $[r_i]_{i=1}^{k-1} \in \mathbb{F}^{k-1}$, a NARK proof

$$\pi.x = [C_i]_{i=1}^k, \pi.\mathbf{w} = [\mathbf{m}_i]_{i=1}^k$$

where $[C_i]_{i=1}^k \in \mathcal{C}^k$ are commitments and $[\mathbf{m}_i]_{i=1}^k$ are prover messages in the special-sound protocol Π_{sps} , and a slack variable μ , the predicate checks that (i) $r_i = \rho_{\text{NARK}}(r_{i-1}, C_i)$ for all $i \in [k-1]$ (where $r_0 := \rho_{\text{NARK}}(\text{pi})$), (ii) $\text{Commit}(\text{ck}, \mathbf{m}_i) = C_i$ for all $i \in [k]$, and (iii)

$$V_{\text{sps}}(\text{pi}, \pi.x, \pi.\mathbf{w}, [r_i]_{i=1}^{k-1}, \mu) := \sum_{j=0}^d \mu^{d-j} \cdot f_j^{\text{V}_{\text{sps}}}(\text{pi}, \pi.\mathbf{w}, [r_i]_{i=1}^{k-1}) = \mathbf{e}$$

where $\mathbf{e} = \mathbf{0}^\ell$ and $\mu = 1$ for the NARK verifier V_{NARK} . Here $f_j^{\text{V}_{\text{sps}}}$ is a degree- j homogeneous algebraic map that outputs ℓ field elements. Degree- j homogeneity says that each monomial term of $f_j^{\text{V}_{\text{sps}}}$ has degree exactly j .

Remark 1. Without loss of generality, we assume that the public input pi is of constant size, as otherwise, we can set it as the hash of the original public input.

Accumulator. The accumulator has the following format:

- *Accumulator instance* $\text{acc}.x := \{\text{pi}, [C_i]_{i=1}^k, [r_i]_{i=1}^{k-1}, E, \mu\}$, where $\text{pi} \in \mathcal{M}^{\ell_{\text{in}}}$ is the accumulated public input, $[C_i]_{i=1}^k \in \mathcal{C}^k$ are the accumulated commitments, $[r_i]_{i=1}^{k-1} \in \mathbb{F}^{k-1}$ are the accumulated challenges, $E \in \mathcal{C}$ is the accumulated commitment to the error terms, and $\mu \in \mathbb{F}$ is a slack variable.
- *Accumulator witness* $\text{acc}.w := \{[\mathbf{m}_i]_{i=1}^k\}$, where $[\mathbf{m}_i]_{i=1}^k$ are the accumulated prover messages.

Accumulation prover. On input commitment key ck (which can be hardwired in the prover's algorithm), accumulator acc , an instance-proof pair (pi, π) where

$$\begin{aligned} \text{acc} &:= (\text{acc}.x = \{\text{pi}', [C'_i]_{i=1}^k, [r'_i]_{i=1}^{k-1}, E, \mu\}, \text{acc}.w = \{[\mathbf{m}'_i]_{i=1}^k\}), \\ \pi &:= (\pi.x = [C_i]_{i=1}^k, \pi.w = [\mathbf{m}_i]_{i=1}^k), \end{aligned}$$

the accumulation prover P_{acc} works as in Figure 3.

Accumulation verifier. On input public input pi , NARK proof instance $\pi.x$, accumulator instance $\text{acc}.x$, accumulation proof pf , and the updated accumulator instance $\text{acc}' .x := \{\text{pi}'', [C''_i]_{i=1}^k, [r''_i]_{i=1}^k, E', \mu'\}$, the accumulation verifier V_{acc} works as in Figure 4.

Decider. On input the commitment key ck (which can be hardwired) and an accumulator

$$\text{acc} = (\text{acc}.x = \{\text{pi}, [C_i]_{i=1}^k, [r_i]_{i=1}^{k-1}, E, \mu\}, \text{acc}.w = \{[\mathbf{m}_i]_{i=1}^k\}),$$

the decider does the checks described in Figure 5.

Remark 2. The accumulation scheme for V_{NARK} is also naturally a folding scheme as defined in Nova [KST22], where we can view an accumulator as a relaxed NP instance with error terms. A NARK proof π is an accumulator with $\mu = 1$ and $E = 0 \in \mathbb{G}$. We can use the same accumulation scheme to fold two accumulators $(\text{acc}, \text{acc}')$ into a new accumulator acc'' . The scheme is identical to the one presented above but with non-trivial μ, \mathbf{e}, E terms for acc . The verifier performs one additional group scalar multiplication. In the language of folding schemes, we can fold two NARK instances into an accumulator; or fold a NARK instance and an accumulator into an updated accumulator; or fold two accumulators into an updated accumulator.

$\mathbf{P}_{\text{acc}}^{\rho_{\text{acc}}, \rho_{\text{NARK}}}(\text{ck}, \text{acc}, (\text{pi}, \pi))$

1. $r_i \leftarrow \rho_{\text{NARK}}(r_{i-1}, C_i) \forall i \in [k-1]$ where $r_0 := \rho_{\text{NARK}}(\text{pi})$.
2. Compute $[\mathbf{e}_j]_{j=1}^{d-1} \in (\mathbb{F}^\ell)^{d-1}$, such that

$$\begin{aligned} & \sum_{j=0}^d (X + \mu)^{d-j} \cdot f_j^{\text{V}_{\text{sps}}} (X \cdot \text{pi} + \text{pi}', [X \cdot \mathbf{m}_i + \mathbf{m}'_{i=1}]^k, [X \cdot r_i + r'_{i=1}]^{k-1}) \\ &= \sum_{j=0}^d \mu^{d-j} f_j^{\text{V}_{\text{sps}}} (\text{pi}', [\mathbf{m}'_{i=1}]^k, [r'_{i=1}]^{k-1}) + X^d \cdot \text{V}_{\text{NARK}}(\text{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}) + \sum_{j=1}^{d-1} \mathbf{e}_j X^j \\ &= \mathbf{e} + \sum_{j=1}^{d-1} \mathbf{e}_j X^j \end{aligned}$$

3. $E_j \leftarrow \text{Commit}(\text{ck}, \mathbf{e}_j) \forall j \in [d-1]$
4. $\alpha \leftarrow \rho_{\text{acc}}(\text{acc}.x, \text{pi}, \pi.x, [E_j]_{j=1}^{d-1}) \in \mathbb{F}$
5. Set vectors

$$\mathbf{v} := \left(1, \text{pi}, [r_i]_{i=1}^{k-1}, [C_i]_{i=1}^k, [\mathbf{m}_i]_{i=1}^k \right), \mathbf{v}' := \left(\mu, \text{pi}', [r'_i]_{i=1}^{k-1}, [C'_i]_{i=1}^k, [\mathbf{m}'_i]_{i=1}^k \right).$$

6. $\mathbf{v}'' := \left(\mu', \text{pi}'', [r''_i]_{i=1}^{k-1}, [C''_i]_{i=1}^k, [\mathbf{m}''_i]_{i=1}^k \right) \leftarrow \alpha \cdot \mathbf{v} + \mathbf{v}'$.
7. $E' \leftarrow E + \sum_{j=1}^{d-1} \alpha^j \cdot E_j$.
8. Set $\text{acc}'.x := \{\text{pi}'', [C''_i]_{i=1}^k, [r''_i]_{i=1}^{k-1}, E', \mu'\}$, $\text{acc}'.\mathbf{w} := \{[\mathbf{m}''_i]_{i=1}^k\}$.
9. Set accumulation proof $\text{pf} := [E_j]_{j=1}^{d-1}$

Figure 3: Accumulation Prover for low-degree Fiat-Shamired NARKs

$\mathbf{V}_{\text{acc}}^{\rho_{\text{acc}}, \rho_{\text{NARK}}}(\text{pi}, \pi.x = [C_i]_{i=1}^k, \text{acc}.x = (\text{pi}', [C'_i]_{i=1}^k, [r'_i]_{i=1}^{k-1}, E, \mu), \text{pf} = [E_j]_{j=1}^{d-1}, \text{acc}'.x)$

1. $r_i \leftarrow \rho_{\text{NARK}}(r_{i-1}, C_i) \forall i \in [k-1]$ where $r_0 := \rho_{\text{NARK}}(\text{pi})$.
2. $\alpha \leftarrow \rho_{\text{acc}}(\text{acc}.x, \text{pi}, \pi.x, \text{pf})$
3. Set vectors

$$\mathbf{v} := \left(1, \text{pi}, [r_i]_{i=1}^{k-1}, [C_i]_{i=1}^k \right), \mathbf{v}' := \text{acc}.x. \left(\mu, \text{pi}', [r'_i]_{i=1}^{k-1}, [C'_i]_{i=1}^k \right).$$

4. Check $\text{acc}'.x. \left(\mu', \text{pi}'', [r''_i]_{i=1}^{k-1}, [C''_i]_{i=1}^k \right) \stackrel{?}{=} \alpha \cdot \mathbf{v} + \mathbf{v}'$.
5. Check $\text{acc}'.x.E' \stackrel{?}{=} \text{acc}.x.E + \sum_{j=1}^{d-1} \alpha^j \cdot E_j$.

Figure 4: Accumulation Verifier for low-degree Fiat-Shamired NARKs

$$D_{\text{acc}}(\text{acc} = (\text{acc.x} = \{\mathbf{pi}, [C_i]_{i=1}^k, [r_i]_{i=1}^{k-1}, E, \mu\}, \text{acc.w} = \{[\mathbf{m}_i]_{i=1}^k\}))$$

1. $C_i \stackrel{?}{=} \text{Commit}(\text{ck}, \mathbf{m}_i)$ for all $i \in [k]$.
2. $\mathbf{e} \leftarrow \sum_{j=0}^d \mu^{d-j} f_j^{\text{V}_{\text{sps}}}(\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1})$ where $f_j^{\text{V}_{\text{sps}}}$ is the degree- j homogeneous algebraic map described in the accumulated predicate.
3. $E \stackrel{?}{=} \text{Commit}(\text{ck}, \mathbf{e})$.

Figure 5: Accumulation Decider for low-degree Fiat-Shamired NARKs

Complexity. Let Π_{sps} be a $(2k-1)$ -move special-sound protocol with the verifier checking ℓ degree- d equations. Denote by $|M|$ the number of elements in prover messages and $|M^*|$ the number of non-zero elements in the prover messages. Assume that \mathbf{pi} is a hash with length 1 (this saves the call $r_0 := \rho_{\text{NARK}}(\mathbf{pi})$), and let $|R|$ be the number of elements in verifier's challenges. We analyze the computational complexity of the accumulation scheme:

- The *accumulation prover*
 - asks $k-1$ queries to ρ_{NARK} and 1 query to ρ_{acc} ;
 - computes $E_j = \text{Commit}(\text{ck}, \mathbf{e}_j)$ for all $j \in [d-1]$, where $\mathbf{e}_j \in \mathbb{F}^\ell$;
 - performs $|R| + |M^*| + 2$ \mathbb{F} -ops to combine $(\mu, \mathbf{pi}, [r_i]_{i=1}^{k-1}, [\mathbf{m}_i]_{i=1}^k)$;
 - performs k \mathbb{G} -ops to combine $[C_i]_{i=1}^k$;
 - computes the coefficients of ℓ degree- d polynomials for $[\mathbf{e}_j]_{j=1}^{d-1}$.
- The *accumulation verifier* performs
 - asks $k-1$ constant size queries to ρ_{NARK} and 1 d -sized query to ρ_{acc} ;
 - $|R| + 2$ \mathbb{F} -ops to combine $(\mu, \mathbf{pi}, [r_i]_{i=1}^{k-1})$;
 - k \mathbb{G} -ops to combine $[C_i]_{i=1}^k$;
 - $d-1$ \mathbb{G} -ops to add $[E_j]_{j=1}^{d-1}$ onto E .
- The *decider*
 - computes $C_i = \text{Commit}(\text{ck}, \mathbf{m}_i)$ for $i \in [k]$ and $E = \text{Commit}(\text{ck}, \mathbf{e})$, with total complexity around $|M| + \ell$ \mathbb{G} -ops.
 - evaluate ℓ degree- d multivariate polynomials to compute vector \mathbf{e} .

Theorem 2. Let $(P_{\text{NARK}}, V_{\text{NARK}})$ be the RO-NARK defined in Section 3.3. Let $\text{cm} = (\text{Setup}, \text{Commit})$ be a binding, homomorphic commitment scheme. Let ρ_{acc} be another random oracle. The accumulation scheme $(P_{\text{acc}}, V_{\text{acc}}, D_{\text{acc}})$ for V_{NARK} satisfies perfect completeness and has knowledge error $(Q+1) \frac{d+1}{|\mathbb{F}|} + \text{negl}(\lambda)$ as defined in Definition 8, against any randomized polynomial-time Q -query adversary.

Proof.

Completeness. Consider any tuple $((\mathbf{pi}, \pi), \text{acc}) \in \mathcal{R}_{\text{acc}}$, that is, $V_{\text{NARK}}(\mathbf{pi}, \pi)$ and $D(\text{acc})$ both accept. Let $(\text{acc}', \mathbf{pf})$ denote the output of the accumulation prover $P_{\text{acc}}(\text{ck}, \text{acc}, (\mathbf{pi}, \pi))$. We argue that both the decider $D(\text{acc}')$ and the accumulation verifier $V_{\text{acc}}(\mathbf{pi}, \pi.x, \text{acc}.x, \mathbf{pf}, \text{acc}'.x)$ will accept, which finishes the proof of perfect completeness by Definition 8.

V_{acc} accepts as P_{acc} and V_{acc} go through the same process of computing challenges $[r_i]_{i=1}^{k-1}$ and α , thus the linear combinations of $\text{acc}.x$ and $(\mathbf{pi}, \pi.x; \mathbf{pf}, [r_i]_{i=1}^{k-1})$ via α will be consistent.

We prove that $D(\text{acc}')$ accepts by scrutinizing the following decider checks.

The check $\text{acc}'.C_i \stackrel{?}{=} \text{Commit}(\text{ck}, \text{acc}'.\mathbf{m}_i)$ succeeds for all $i \in [k]$. This is because

$$\text{acc}'.\{C_i, \mathbf{m}_i\} = \text{acc}.\{C_i, \mathbf{m}_i\} + \alpha \cdot \pi.\{C_i, \mathbf{m}_i\}$$

for all $i \in [k]$, where $\pi.C_i = \text{Commit}(\text{ck}, \pi.\mathbf{m}_i)$ because $V_{\text{NARK}}(\mathbf{pi}, \pi)$ accepts, and $\text{acc}.C_i = \text{Commit}(\text{ck}, \text{acc}.\mathbf{m}_i)$ because $D(\text{acc})$ accepts. Thus the check succeeds by the homomorphism of the commitment scheme.

The decider computes $\mathbf{e}' \leftarrow \sum_{j=0}^d (\text{acc}'.\mu)^{d-j} f_j^{\text{V}_{\text{sps}}}(\text{acc}'.\{\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}\})$ such that for $\mathbf{e} = \sum_{j=0}^d \text{acc}.\mu^{(d-j)} \cdot f_j^{\text{V}_{\text{sps}}}(\text{acc}.\{\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}\})$, it holds that

$$\begin{aligned} \mathbf{e}' &= \mathbf{e} + \sum_{j=1}^{d-1} \alpha^j \cdot \mathbf{pf}.\mathbf{e}_j \\ &= \sum_{j=0}^d (\alpha + \text{acc}.\mu)^{d-j} \cdot f_j^{\text{V}_{\text{sps}}}(\alpha \cdot \{\mathbf{pi}, \pi.[\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}\} + \text{acc}.\{\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}\}). \end{aligned}$$

By the definition of $\mathbf{pf}.\mathbf{e}_j$ and the homomorphism of the commitment scheme, and because $D(\text{acc})$ accepts and checks $E = \text{Commit}(\text{ck}, \mathbf{e})$, we have that $E' = \text{Commit}(\text{ck}, \mathbf{e}')$.

Knowledge-Soundness. We show that the scheme has knowledge-soundness by showing that there exists an underlying $(d+1)$ -special-sound protocol and then applying the Fiat-Shamir transform to show that the accumulation scheme is knowledge sound. Consider the public-coin interactive protocol $\Pi_I = (P_I(\mathbf{pi}, \pi, \text{acc}), V_I(\mathbf{pi}, \pi.x, \text{acc}.x))$ where P_I sends $\mathbf{pf} = [E_j]_{j=1}^{d-1} \in \mathbb{G}^{d-1}$ as computed by P_{acc} to V_I . The verifier sends a random challenge $\alpha \in \mathbb{F}$, and the prover P_I responds with acc' as computed by P_{acc} . V_I accepts if $D_{\text{acc}}(\text{acc}') = 0$ and $V_{\text{acc}}(\mathbf{pi}, \pi.x, \text{acc}.x, \mathbf{pf}, \text{acc}'.x) = 0$ using the random challenge α , instead of a Fiat-shamir challenge.

Claim 1: Π_I is $(d+1)$ -special-sound Consider the relation \mathcal{R}_{acc} where \mathcal{R}_{acc} is defined in Definition 8. Consider $d+1$ accepting transcripts for Π_I :

$$\{\mathcal{T}_i := (\mathbf{pi}, \pi.x, \text{acc}.x; \text{acc}'_i, \mathbf{pf}_i)\}_{i=1}^{d+1}.$$

We construct an extractor Ext_{acc} that extracts a witness for $\mathcal{R}_{\text{acc}}(\text{pi}, \pi.x, \text{acc}.x)$ given \mathcal{T} .

For all $i \in [d+1]$,

$$(\text{acc}'_i) = (\mu'_i, \text{pi}'_i, [C'_{i,j}]_{j=1}^k, [r_{i,j}]_{j=1}^{k-1}, E'_i, [\mathbf{m}'_{i,j}]_{j=1}^k)$$

and $\text{pf}_i = \text{pf} = [E_j]_{j=1}^{d-1}$.

Given that the transcripts are accepting, i.e. both V_{acc} and D_{acc} accept, we have that $\text{Commit}(\text{ck}, \mathbf{e}'_i) = E'_i = \text{acc}.E + \sum_{j=1}^{d-1} \alpha_i^j E_j$ for all $i \in [d+1]$, whereas

$$\mathbf{e}'_i := \sum_{j=0}^d \mu_i^{d-j} f_j^{\mathcal{R}}(\pi'_i, [\mathbf{m}'_{i,j}]_{j=1}^k, [r_{i,j}]_{j=1}^{k-1}).$$

Using a Vandermonde matrix of the challenges $\alpha_1, \dots, \alpha_d$ we can compute $\mathbf{e}, [\mathbf{e}_j]_{j=1}^{d-1}$ such that $E_j = \text{Commit}(\text{ck}, \mathbf{e}_j)$ and $\text{acc}.E = \text{Commit}(\text{ck}, \mathbf{e})$ from the equations above. Therefore we have that $\mathbf{e}'_i = \mathbf{e} + \sum_{j=1}^{d-1} \alpha_i^j \mathbf{e}_j$ for all $i \in [d+1]$.

Additionally using two challenges (α_1, α_2) , Ext_{acc} can compute $\pi.\mathbf{w} = [\mathbf{m}_j]_{j=1}^k = [\frac{\text{acc}.\mathbf{m}_{1,j} - \text{acc}'.\mathbf{m}_{2,j}}{\alpha_1 - \alpha_2}]_{j=1}^k$. It holds that $\text{acc}.\mathbf{m}_j = \text{acc}'.\mathbf{m}_{1,j} - \alpha_1 \cdot \pi.\mathbf{m}_j \forall j \in [k]$, such that $\pi.C_j = \text{Commit}(\text{ck}, \pi.\mathbf{m}_j)$ and $\text{acc}.C_j = \text{Commit}(\text{ck}, \text{acc}.\mathbf{m}_j)$. If for any other challenge and any j , $\text{acc}'.\mathbf{m}_j \neq \alpha \pi.\mathbf{m}_j + \text{acc}.\mathbf{m}_j$, then this can be used to compute a break of the commitment scheme cm . This happens with negligible probability by assumption.

Otherwise, we have that $\sum_{j=0}^d \mu_i^{d-j} f_j^{\mathcal{R}}(\pi_j, [\mathbf{m}_{i,j}]_{i=1}^k, [r_{i,j}]_{i=1}^{k-1}) - \mathbf{e}_i = 0$ for all $i \in [d+1]$. Together this implies that the degree d polynomial

$$\begin{aligned} p(X) &= \sum_{j=0}^d (X + \text{acc}.\mu)^{d-j} \cdot f_j^{\text{V}_{\text{sps}}} (X \cdot \text{pi} + \text{acc}.\text{pi}, [X \cdot \mathbf{m}_i + \text{acc}.\mathbf{m}_i]_{i=1}^k, [X \cdot r_i + \text{acc}.r_i]_{i=1}^{k-1}) \\ &\quad - \mathbf{e} - \sum_{j=1}^{d-1} \mathbf{e}_j X^j, \end{aligned} \tag{2}$$

is zero on $d+1$ points $(\alpha_1, \dots, \alpha_{d+1})$, i.e. is zero everywhere. The constant term of this polynomial is

$$\sum_{j=0}^d \text{acc}.\mu^{d-j} \cdot f_j^{\text{V}_{\text{sps}}} (\text{acc}.\text{pi}, [\text{acc}.\mathbf{m}_i]_{i=1}^k, [\text{acc}.r_i]_{i=1}^{k-1}) - \mathbf{e}.$$

It being 0 implies that $D(\text{acc}) = 0$. Additionally, the degree d term of the polynomial is

$$\sum_{j=0}^d f_j^{\text{V}_{\text{sps}}} (\text{pi}, [\pi.\mathbf{m}_i]_{i=1}^k, [\pi.r_i]_{i=1}^{k-1}).$$

Together with V_{acc} checking that the challenges r_i are computed correctly this implies that $V_{\text{NARK}}(\text{pi}, \pi) = 0$. Ext thus outputs a valid witness $(\pi.\mathbf{w}, \text{acc}.\mathbf{w}) \in \mathcal{R}_{\text{acc}}(\text{pi}, \pi.x, \text{acc}.x)$ and

thus Π_I is $(d+1)$ -special-sound. Using Lemma 1, we have that $\Pi_{AS} = \text{FS}[\Pi_I]$ is a NARK for \mathcal{R}_{acc} with knowledge soundness $(Q+1) \cdot \frac{d+1}{|\mathbb{F}|} + \text{negl}(\lambda)$. This implies that acc is an accumulation scheme with $((Q+1) \cdot \frac{d+1}{|\mathbb{F}|} + \text{negl}(\lambda))$ -knowledge soundness. \square

3.5 Compressing verification checks for high-degree verifiers

Observe that the accumulation prover needs to perform $\Omega(d\ell)$ group operations to commit to the $d-1$ error vectors $\mathbf{e}_j \in \mathbb{F}^\ell$ ($1 \leq j < d$); and the accumulation verifier needs to check the combination of d error vector commitments. This can be a bottleneck when the verifier degree d is high. In this circumstance, we can optimize the accumulation complexity by transforming the underlying special-sound protocol Π_{sps} into a new special-sound protocol $\text{CV}[\Pi_{\text{sps}}]$ for the same relation \mathcal{R} . This optimization compresses the ℓ degree- d equations checked by the verifier into a single degree- $(d+2)$ equation using a random linear combination, with the tradeoff of additionally checking $2\sqrt{\ell}$ degree-2 equations. We describe the generic transformation below.

Compressing verification checks. W.l.o.g. assume ℓ is a perfect square, then we can transform Π_{sps} into a special-sound protocol $\text{CV}[\Pi_{\text{sps}}]$ where the \mathbf{V}_{sps} reduces from ℓ degree- d checks to 1 degree- $(d+2)$ check and additionally $2\sqrt{\ell}$ degree-2 checks. Instead of checking the output of \mathbf{V}_{sps} to be ℓ zeroes, we take a random linear combination of the ℓ verification equations using powers of a challenge β . For example, if the map is $\mathbf{V}_{\text{sps}}(x_1, x_2) := (f_1(x_1, x_2), f_2(x_1, x_2)) = (x_1 + x_2, x_1x_2)$ we can set the new algebraic map as $\mathbf{V}'_{\text{sps}}(x_1, x_2, \beta) := f_1(x_1, x_2) + \beta \cdot f_2(x_1, x_2) = (x_1 + x_2) + \beta x_1x_2$ for a random β . Doing this naively reduces the output length to 1 but also requires the verifier to compute the appropriate powers of β . This would increase the degree by ℓ , an undesirable tradeoff. To mitigate this, we can have the prover precompute powers of β , i.e. $\beta, \beta^2, \dots, \beta^\ell$ and send them to the verifier. The verifier then only needs to check consistency between the powers of β , which can be done using a degree 2 check, e.g. $\beta^{i+1} = \beta^i \cdot \beta$ and the degree d verification equation increases in degree by 1. This mitigates the degree increase but requires the prover to send another message of length ℓ . To achieve a more optimal tradeoff, we write each $i = j + k \cdot \sqrt{\ell}$ for $j, k \in [1, \sqrt{\ell}]$. The prover then sends $\sqrt{\ell}$ powers of β and $\sqrt{\ell} - 1$ powers of $\beta^{\sqrt{\ell}}$. From these, each power of β from 1 to ℓ can be recomputed using just one multiplication. This results in the prover sending an additional message of length $2\sqrt{\ell}$, the original ℓ verification checks being transformed into a single degree $d+2$ check and additionally $2\sqrt{\ell}$ degree 2 checks for the consistency of the powers of β .

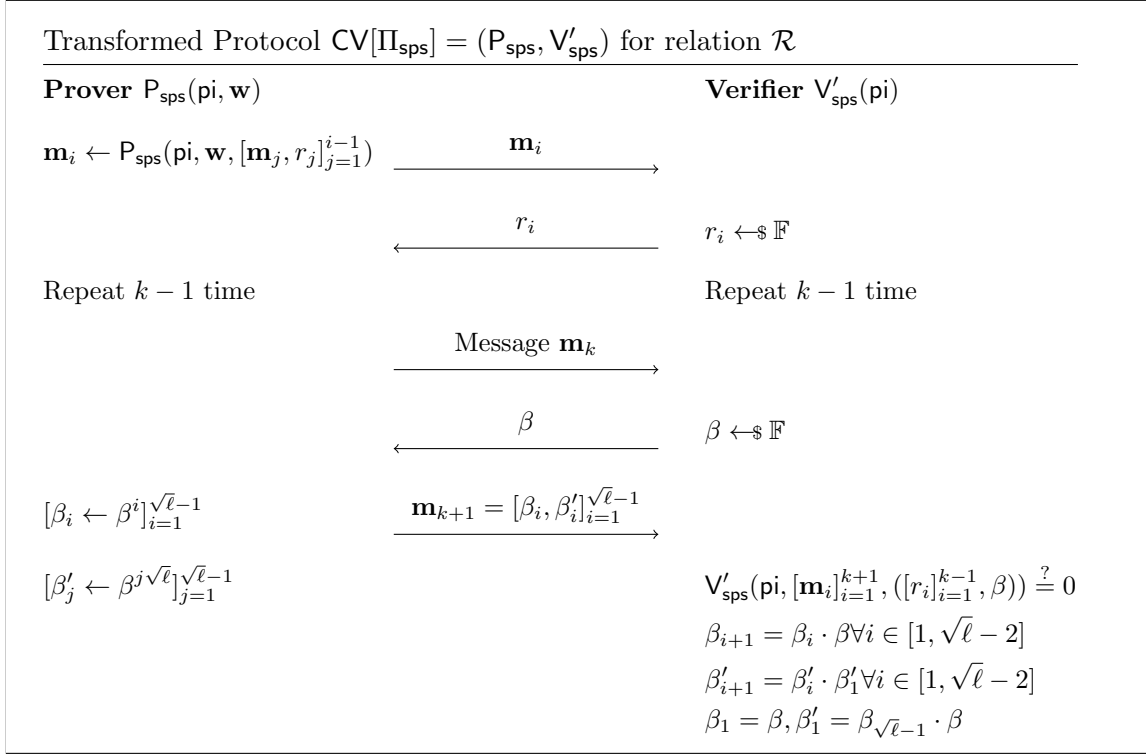


Figure 6: Compressed verification of Π_{sps} .

We describe the transformed protocol in Figure 6, where

$$\begin{aligned} V'_{\text{sps}}(\mathbf{pi}, [\mathbf{m}_i]_{i=1}^{k+1}, ([r_i]_{i=1}^{k-1}, \beta)) &:= \sum_{i=0}^{\sqrt{\ell}-1} \sum_{j=0}^{\sqrt{\ell}-1} \beta_i \cdot \beta'_j \cdot V_{\text{sps}, i+j\sqrt{\ell}}(\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}) \\ &= \sum_{j=0}^{\ell-1} \beta^j \cdot V_{\text{sps}, j}(\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}) \end{aligned}$$

and $V_{\text{sps}, j}(\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1})$ is the $(j+1)$ -th ($0 \leq j < \ell$) equation checked by V_{sps} . The transformed protocol is a $(2k+1)$ -move special-sound protocol for the same relation \mathcal{R} . The transformed verifier now checks 1 degree- $(d+2)$ equation and additionally $2\sqrt{\ell}$ degree-2 equations.

Lemma 3. *Let Π_{sps} be a $(2k-1)$ -move protocol for relation \mathcal{R} with (a_1, \dots, a_{k-1}) -special-soundness, in which the verifier outputs ℓ elements. The transformed protocol $\text{CV}[\Pi_{\text{sps}}]$ of Π_{sps} is $(a_1, \dots, a_{k-1}, \ell)$ -special-sound.*

Proof. Let Ext_{sps} be the extractor for Π_{sps} . We construct an extractor Ext_{CV} of $\text{CV}[\Pi_{\text{sps}}]$ for the same relation \mathcal{R} . Given an $(a_1, \dots, a_{k-1}, \ell)$ -tree \mathcal{T} of accepting transcripts, Ext_{CV} invokes Ext_{sps} on input the depth- $(k-1)$ transcript subtree of \mathcal{T} , and return what Ext_{sps} outputs.

We prove that the extractor succeeds. For each internal node u at depth $k-1$, it has ℓ children where each child maps to a distinct value of $\beta \in \mathbb{F}$. Fix the messages $\text{msg} = (\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1})$ at node u and let $V_{\text{sps}} := (V_{\text{sps}, 1}, \dots, V_{\text{sps}, \ell})$ be the verifier of Π_{sps} . Define the degree $\ell-1$ univariate polynomial

$$p(X) := \sum_{j=0}^{\ell-1} X^j \cdot c_j$$

where $c_j := V_{\text{sps}, j}(\text{msg}) \in \mathbb{F}$ is $V_{\text{sps}, j}$'s output on message msg . Since the transcripts are accepting, it holds that p evaluates to zero on the ℓ different values of β that correspond to the ℓ children of node u . Thus the univariate polynomial p is a zero polynomial, which implies that V_{sps} outputs zero vector on message msg . Therefore for every node u at depth $k-1$, the sub-transcript from root to node u is an accepting transcript to Π_{sps} . Therefore the input to Ext_{sps} is a valid (a_1, \dots, a_{k-1}) -tree of accepting transcripts, and Ext_{sps} will output the correct witness. \square

High-low degree accumulation. After the transformation, the error vectors \mathbf{e}_j ($1 \leq j \leq d+1$) become single field elements, and we can use the trivial commitment $E_j := \text{Commit}(\text{ck}, e_j) := e_j$ without group operations. Additionally, we can use a separate error vector $\mathbf{e}' \in \mathbb{F}^{2\sqrt{\ell}}$ to keep track of the error terms for the $2\sqrt{\ell}$ degree-2 checks, and set

$E' := \text{Commit}(\text{ck}, \mathbf{e}') \in \mathbb{G}$ to be the corresponding error commitment. The accumulation prover only needs to perform $O(\sqrt{\ell})$ additional group operations to commit \mathbf{m}_{k+1} and \mathbf{e}' , and compute the coefficients of a degree- $(d+2)$ univariate polynomial, which is described as the sum of $O(\ell)$ polynomials. The accumulator instance needs to include one more challenge β and two commitments (for \mathbf{m}_{k+1} and \mathbf{e}'). The accumulator verifier needs to do only $k+2$ (rather than $k+d-1$) group scalar multiplications, with the tradeoff of 1 more hash and $O(d)$ more field operations. This high-low degree accumulation is described in detail in Appendix A.

Theorem 3 (IVC for high-degree special-sound protocols). *Let \mathbb{F} be a finite field, such that $|\mathbb{F}| \geq 2^\lambda$ and $\text{cm} = (\text{Setup}, \text{Commit})$ be a binding homomorphic commitment scheme for vectors in \mathbb{F} . Let $\Pi_{\text{sps}} = (\text{P}_{\text{sps}}, \text{V}_{\text{sps}})$ be a special-sound protocol for an NP-complete relation \mathcal{R}_{NP} with the following properties:*

- It's $(2k-1)$ move.
- It's (a_1, \dots, a_{k-1}) -out-of- $|\mathbb{F}|$ special-sound. Such that the knowledge error $\kappa = 1 - \prod_{i=1}^{k-1} (1 - \frac{a_i}{|\mathbb{F}|}) = \text{negl}(\lambda)$
- The inputs are in $\mathbb{F}^{\ell_{\text{in}}}$
- The verifier is degree $d = \text{poly}(\lambda)$ with output in \mathbb{F}^ℓ

Then, under the Fiat-Shamir heuristic for a cryptographic hash function H (Definition 9), there exist two IVC schemes $\text{IVC} = (\text{P}_{\text{IVC}}, \text{V}_{\text{IVC}})$ and $\text{IVC}_{\text{CV}} = (\text{P}_{\text{CV,IVC}}, \text{V}_{\text{CV,IVC}})$ with predicates expressed in \mathcal{R}_{NP} with the following efficiencies:

	No CV	CV
$\text{P}_{\text{IVC}} \text{ native}$	$\sum_{i=1}^k \mathbf{m}_i^* + (d-1)\ell\mathbb{G}$ $\text{P}_{\text{sps}} + L(\text{V}_{\text{sps}}, d)$	$\sum_{i=1}^k \mathbf{m}_i^* + O(\sqrt{\ell})\mathbb{G}$ $\text{P}_{\text{sps}} + L'(\text{V}_{\text{sps}}, d+2)$
$\text{P}_{\text{IVC}} \text{ recursive}$	$k+d-1\mathbb{G}$ $k + \ell_{\text{in}}\mathbb{F}$ $(k+d+O(1))H + 1\text{H}_{\text{in}}$	$k+2\mathbb{G}$ $k + \ell_{\text{in}} + d + 1\mathbb{F}$ $(k+d+O(1))H + 1\text{H}_{\text{in}}$
$\text{V}_{\text{IVC}}:$	$\ell + \sum_{i=1}^k \mathbf{m}_i \mathbb{G}$ V_{sps}	$O(\sqrt{\ell}) + \sum_{i=1}^k \mathbf{m}_i \mathbb{G}$ $O(\ell) + \text{V}_{\text{sps}}$
$ \pi_{\text{IVC}} :$	$k + \ell_{\text{in}}\mathbb{F}$ $k + 1\mathbb{G}$ $\sum_{i=1}^k \mathbf{m}_i $	$k + \ell_{\text{in}} + 1\mathbb{F}$ $k + 2\mathbb{G}$ $\sum_{i=1}^k \mathbf{m}_i + O(\sqrt{\ell})$

The first row displays the native operations of the IVC prover. The second row describes the size of the recursive statement expressed as an instance of \mathcal{R}_{NP} for which P_{IVC} creates a proof. The third row is the computation of V_{IVC} , and the last row is the size of the proof.

In the table, $|\mathbf{m}_i|$ denotes the prover message length; $|\mathbf{m}_i^*|$ is the number of non-zero elements in \mathbf{m}_i ; \mathbb{G} for rows 1-3 is the total length of the messages committed using Commit .

\mathbb{F} are field operations. H are calls to the random oracle and H_{in} are the hash to the public input and accumulator instance. P_{sps} (and V_{sps}) is the cost of running the prover (and the algebraic verifier) of the special-sound protocol, respectively. $L(V_{\text{sps}}, d)$ is the cost of computing the coefficients of the degree d polynomial

$$e(X) := \sum_{j=0}^d (\mu + X)^{d-j} \cdot f_j^{V_{\text{sps}}}(\text{acc} + X \cdot \pi), \quad (3)$$

and $L'(V_{\text{sps}}, d+2)$ is the cost of computing the coefficients of the degree $d+2$ polynomial

$$e(X) := \sum_{a=0}^{\sqrt{\ell}-1} \sum_{b=0}^{\sqrt{\ell}-1} (X \cdot \pi \cdot \beta_a + \text{acc} \cdot \beta_a)(X \cdot \pi \cdot \beta'_b + \text{acc} \cdot \beta'_b) \sum_{j=0}^d (\mu + X)^{d-j} \cdot f_{j, a+b\sqrt{\ell}}^{V_{\text{sps}}}(\text{acc} + X \cdot \pi), \quad (4)$$

where all inputs are linear functions in a formal variable X , and $f_{j,i}^{V_{\text{sps}}}$ is the i th ($0 \leq i \leq \ell-1$) component of $f_j^{V_{\text{sps}}}$'s output. For the proof size, \mathbb{G} and \mathbb{F} are the number of commitments and field elements, respectively.

Proof. The construction first defines the two NARKs

$$\Pi_{\text{NARK}} = (P_{\text{NARK}}, V_{\text{NARK}}) = \text{FS}[\text{cm}[\Pi_{\text{sps}}]],$$

and

$$\Pi_{\text{NARK,CV}} = (P_{\text{NARK,CV}}, V_{\text{NARK,CV}}) = \text{FS}[\text{cm}[\text{CV}[\Pi_{\text{sps}}]]].$$

Then we construct the accumulation scheme $(P_{\text{acc}}, V_{\text{acc}}) = \text{acc}[\Pi_{\text{NARK}}]$ using the accumulation scheme from Section 3.4 and $(P_{\text{acc,HL}}, V_{\text{acc,HL}}) = \text{acc}_{\text{HL}}[\Pi_{\text{NARK,CV}}]$ using the accumulation scheme from Appendix A. Then we apply the transformation from Theorem 1 to construct the IVC schemes IVC and IVC_{CV}.

Security: By Lemmas 1,2, we have that Π_{NARK} has $(Q+1) \cdot [1 - \prod_{i=1}^{k-1} (1 - \frac{a_i}{|\mathbb{F}|})]$ knowledge error for relation $\mathcal{R}_{\text{cm}}^{\mathcal{R}_{\text{NP}}}$ for a polynomial-time Q -query RO-adversary. Witnesses for $\mathcal{R}_{\text{cm}}^{\mathcal{R}_{\text{NP}}}$ are either a witness for \mathcal{R}_{NP} or a break of the binding property of cm . Assuming that cm is a binding commitment scheme, the probability that a polynomial time adversary and a polynomial time extractor can compute such a break is $\text{negl}(\lambda)$. Thus Π_{NARK} has knowledge error $\kappa = (Q+1) \cdot [1 - \prod_{i=1}^{k-1} (1 - \frac{a_i}{|\mathbb{F}|})] + \text{negl}(\lambda)$ for \mathcal{R}_{NP} . Analogously and using Lemma 3, $\Pi_{\text{NARK,CV}}$ has knowledge soundness with knowledge error $\kappa' = (Q+1) \cdot [1 - (1 - \frac{\ell}{|\mathbb{F}|}) \prod_{i=1}^{k-1} (1 - \frac{a_i}{|\mathbb{F}|})] + \text{negl}(\lambda)$ for \mathcal{R}_{NP} . By assumption, κ and κ' are negligible in λ . Using Theorem 2 and Corollary 2 we can construct accumulation schemes acc and

⁴For example if $f_d = \prod_{i=1}^d (a_i + b_i \cdot X)$ then a naive algorithm takes $O(d^2)$ time but using FFTs it can be computed in time $O(d \log^3 d)$ [CBBZ22].

acc_{CV} for Π_{NARK} and $\Pi_{\text{NARK,CV}}$, respectively. The accumulation schemes have negligible knowledge error as $d = \text{poly}(\lambda)$. Under the Fiat-Shamir heuristic for \mathbf{H} we can turn the NARKs and the accumulation schemes into secure schemes in the standard model.

By Theorem 1, this yields IVC and IVC_{CV} , secure IVC schemes with predicates expressed in \mathcal{R}_{NP} .

Efficiency: We first analyze the efficiency for IVC. The IVC-prover runs P_{sps} to compute all prover messages. It also commits to all the P_{sps} messages using cm . Finally, it needs to compute all error terms $\mathbf{e}_1, \dots, \mathbf{e}_{d-1}$ and commit to them. The error terms are computed by symbolically evaluating the polynomial $e(X)$ in Equation 4 with linear functions as inputs. The recursive circuit combines a new proof $\pi.x$ with an accumulator $\text{acc}.x$. The size of the accumulator instance is ℓ_{in} field elements for the input, $k - 1$ field elements for the interactive-proof challenges, 1 field element for the accumulator challenge, and k commitments for the P_{sps} messages and $d - 1$ commitments for the error terms. The IVC verifier checks the correctness of the commitments and runs V_{sps} .

For IVC_{CV} , the prover needs to additionally commit to a message \mathbf{m}_{k+1} with length $O(\sqrt{\ell})$; the number of error terms also increases from $d - 1$ to $d + 1$. Fortunately, the error terms are only one element in \mathbb{F} , so we can use the identity function as the trivial commitment scheme. Thus, there is no cost for committing to the $d + 1$ error terms when using CV. However, there is another separate error term $\mathbf{e}' \in \mathbb{F}^{2\sqrt{\ell}}$ for the additional $O(\sqrt{\ell})$ degree-2 checks, thus the prover needs to commit to $E' = \text{Commit}(\mathbf{e}')$. The size of the accumulator instance is ℓ_{in} field elements for the input, k field elements for the interactive-proof challenges, 1 field element for the accumulator challenge, $k + 1$ commitments for the prover messages, $d + 1$ field elements for the error terms of the high-degree checks, and 1 commitment for the additional error term \mathbf{e}' . \square

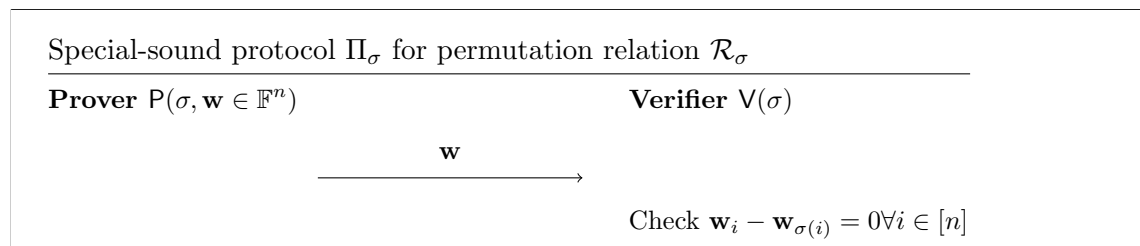
Remark 3. *For simplicity, we assume that the public input, the prover messages, and the verifier challenges are all in the same field \mathbb{F} . This isn't strictly necessary; for example, the challenges could be drawn from a subset of \mathbb{F} . More generally, we can also allow prover messages to be group elements in \mathbb{G} given a homomorphic commitment scheme to group elements (e.g. [AFGHO10]).*

4 Special-sound subprotocols for Protostar

In this section, we present special-sound protocols for permutation, high-degree gate, circuit selection and lookup relations, which are the building blocks for the (non-uniform) Plonkish circuit-satisfiability relations. We can build accumulation schemes for (and thus IVCs from) these special-sound protocols via the framework presented in Section 3.

4.1 Permutation relation

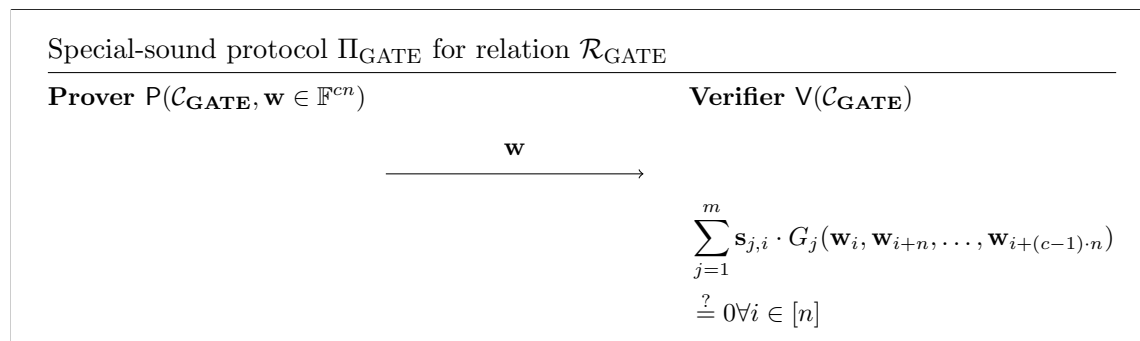
Definition 10. Let $\sigma : [n] \rightarrow [n]$ be a permutation, the relation \mathcal{R}_σ is the set of tuples $\mathbf{w} \in \mathbb{F}^n$ such that $\mathbf{w}_i = \mathbf{w}_{\sigma(i)}$ for all $i \in [n]$.



Complexity. Π_σ is a 1-move protocol (i.e. $k = 1$); the degree of the verifier is 1.

4.2 High-degree custom gate relation

Definition 11. Given configuration $\mathcal{C}_{\text{GATE}} := (n, c, d, [\mathbf{s}_i \in \mathbb{F}^n, G_i]_{i=1}^m)$ where n is the number of gates, c is the arity per gate, d is the gate degree, $[\mathbf{s}_i]_{i=1}^m$ are the selector vectors, and $[G_i]_{i=1}^m$ are the gate formulas, the relation $\mathcal{R}_{\text{GATE}}$ is the set of tuples $\mathbf{w} \in \mathbb{F}^{cn}$ such that $\sum_{j=1}^m \mathbf{s}_{j,i} \cdot G_j(\mathbf{w}_i, \mathbf{w}_{i+n}, \dots, \mathbf{w}_{i+(c-1) \cdot n}) = 0$ for all $i \in [n]$.



Complexity. Π_{GATE} is a 1-move protocol (i.e. $k = 1$) with verifier degree d .

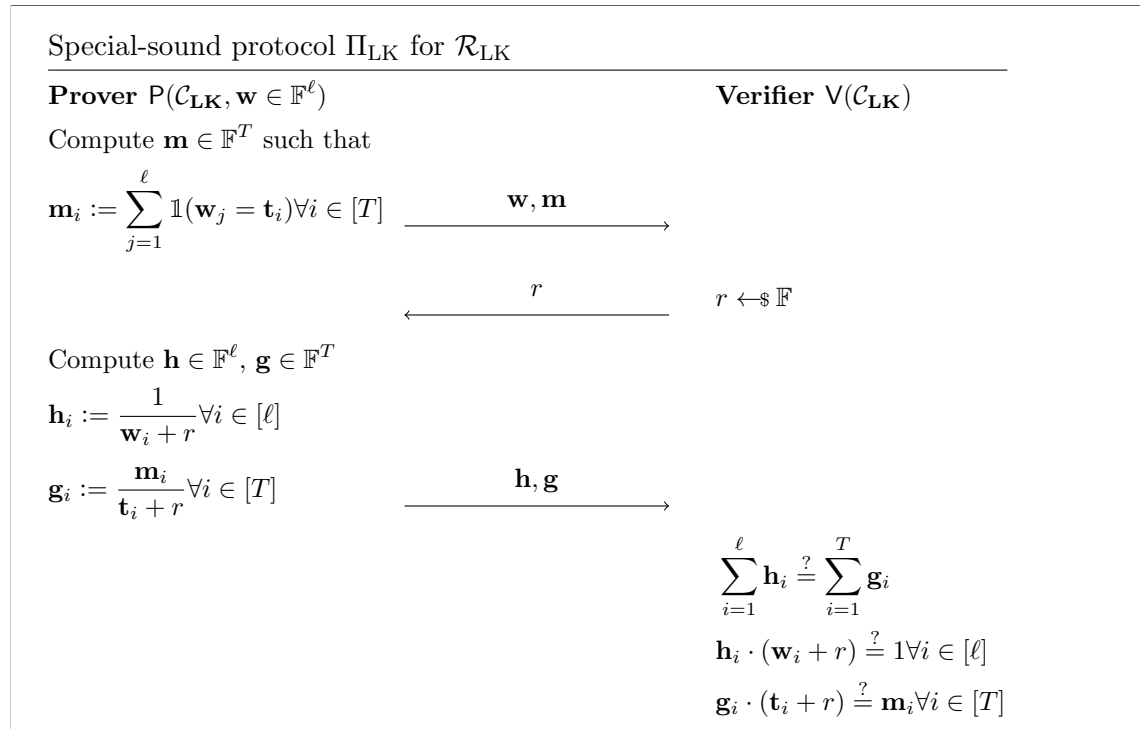
4.3 Lookup relation

Definition 12. Given configuration $\mathcal{C}_{\text{LK}} := (T, \ell, \mathbf{t})$ where ℓ is the number of lookups and $\mathbf{t} \in \mathbb{F}^T$ is the lookup table, the relation \mathcal{R}_{LK} is the set of tuples $\mathbf{w} \in \mathbb{F}^\ell$ such that $\mathbf{w}_i \in \mathbf{t}$ for all $i \in [\ell]$.

We recall a useful lemma for lookup relation from [Hab22], and present a special-sound protocol for the lookup relation.

Lemma 4 (Lemma 5 of [Hab22]). *Let \mathbb{F} be a field of characteristic $p > \max(\ell, T)$. Given two sequences of field elements $[\mathbf{w}_i]_{i=1}^\ell$ and $[\mathbf{t}_i]_{i=1}^T$, we have $\{\mathbf{w}_i\} \subseteq \{\mathbf{t}_i\}$ as sets (with multiples of values removed) if and only if there exists a sequence $[\mathbf{m}_i]_{i=1}^T$ of field elements such that*

$$\sum_{i=1}^{\ell} \frac{1}{X + \mathbf{w}_i} = \sum_{i=1}^T \frac{\mathbf{m}_i}{X + \mathbf{t}_i}. \quad (5)$$



Achieving perfect completeness. Note that the protocol does not have perfect completeness. If there exists an \mathbf{w}_i or \mathbf{t}_i such that $\mathbf{w}_i + r = 0$ or $\mathbf{t}_i + r = 0$ then the prover message is undefined. We can achieve perfect completeness by having the verifier set $\mathbf{h}_i = 0$ or $\mathbf{g}_i = 0$ in this case and changing the verification equations to

$$(\mathbf{w}_i + r) \cdot (\mathbf{h}_i \cdot (\mathbf{w}_i + r) - 1) = 0$$

and

$$(\mathbf{t}_i + r) \cdot (\mathbf{g}_i \cdot (\mathbf{t}_i + r) - \mathbf{m}_i) = 0.$$

These checks ensure that either $\mathbf{h}_i = \frac{1}{\mathbf{w}_i + r}$ or $\mathbf{w}_i + r = 0$. The checks increase the verifier degree to 3. Without these checks, the protocol has a negligible completeness error of $\frac{\ell+T}{|\mathbb{F}|}$. This completeness error can likely be ignored in practice, and these checks do not

need to be implemented. However, to achieve the full definition of PCD (which has perfect completeness) and use Theorem 1 by [BCLMS21], we require that all protocols have perfect completeness.

Complexity. Π_{LK} is a 3-move protocol (i.e. $k = 2$); the degree of the verifier is 2; the number of non-zero elements in the prover message is at most 4ℓ .

Accumulation with $O(\ell)$ prover complexity. The prover complexity of Π_{LK} is due to the sparseness of $\mathbf{g} \in \mathbb{F}^T$ and $\mathbf{m} \in \mathbb{F}^T$. However, there is no guarantee that when building an accumulation scheme for Π_{LK} , the accumulated acc.g and acc.m are sparse. This is an issue, as the prover needs to compute the error term \mathbf{e}_1 . If we expand the accumulation procedures, we see that the three verification checks lead to three components of the error term \mathbf{e}_1 :

$$\mathbf{e}_1^{(1)} = \left(\sum_{i=1}^{\ell} \text{acc.h}_i - \sum_{i=1}^T \text{acc.g}_i \right) + \mu \left(\sum_{i=1}^{\ell} \pi \cdot \mathbf{h}_i - \sum_{i=1}^T \pi \cdot \mathbf{g}_i \right) \in \mathbb{F}$$

$$\mathbf{e}_1^{(2)} = \text{acc.h} \circ (\pi \cdot \mathbf{w} + \pi \cdot r \cdot \mathbf{1}^{\ell}) + \pi \cdot \mathbf{h} \circ (\text{acc.w} + \text{acc.r} \cdot \mathbf{1}^{\ell}) - 2\mu \cdot \mathbf{1}^{\ell} \in \mathbb{F}^{\ell}$$

$$\mathbf{e}_1^{(3)} = \text{acc.g} \circ (\mathbf{t} + \pi \cdot r \cdot \mathbf{1}^T) + \pi \cdot \mathbf{g} \circ (\mu \cdot \mathbf{t} + \text{acc.r} \cdot \mathbf{1}^T) - \mu \cdot \pi \cdot \mathbf{m} - \text{acc.m} \in \mathbb{F}^T.$$

We examine all three components below.

For $\mathbf{e}_1^{(1)}$, we see that $(\sum_{i=1}^{\ell} \pi \cdot \mathbf{h}_i - \sum_{i=1}^T \pi \cdot \mathbf{g}_i) = 0$ by the assumption that π is valid, and $(\sum_{i=1}^{\ell} \text{acc.h}_i - \sum_{i=1}^T \text{acc.g}_i) = \text{acc.e}^{(1)}/\text{acc.}\mu$ (where $\text{acc.e}^{(1)}$ is the first component of the error vector for acc). Thus $\mathbf{e}_1^{(1)} = \text{acc.e}^{(1)}/\text{acc.}\mu$. We observe that since in IVC the accumulator $\text{acc.e}^{(1)}$ is initiated with 0, this implies that for all iterations $\mathbf{e}_1^{(1)} = 0$.

For $\mathbf{e}_1^{(2)}$, it is computed from terms of size ℓ , so can be computed in time $O(\ell)$.

For $\mathbf{e}_1^{(3)}$, note that $\text{acc.}\mu$, acc.r and $\pi \cdot r$ are all scalars. Also note that the accumulation prover only needs to compute the commitment $E_1 = \text{Commit}(\text{ck}, \mathbf{e}_1) = \text{Commit}(\text{ck}, \mathbf{e}_1^{(1)}) + \text{Commit}(\text{ck}, 0 || \mathbf{e}_1^{(2)}) + \text{Commit}(\text{ck}, \mathbf{0}^{\ell+1} || \mathbf{e}_1^{(3)})$, not the actual vector \mathbf{e}_1 . We will compute $E_1^{(3)} = \text{Commit}(\text{ck}, \mathbf{e}_1^{(3)})$ homomorphically from the commitments below (dropping the zero padding for readability):

1. $G = \text{Commit}(\text{ck}, \pi \cdot \mathbf{g})$,
2. $G' = \text{Commit}(\text{ck}, \text{acc.g})$,
3. $M = \text{Commit}(\text{ck}, \pi \cdot \mathbf{m})$,
4. $M' = \text{Commit}(\text{ck}, \text{acc.m})$,
5. $GT = \text{Commit}(\text{ck}, \pi \cdot \mathbf{g} \circ \mathbf{t})$,
6. $GT' = \text{Commit}(\text{ck}, \text{acc.g} \circ \mathbf{t})$.

Given these commitments, we can compute

$$E_1^{(3)} = GT' + \pi \cdot r \cdot G' + \text{acc.}\mu \cdot GT + \text{acc.r} \cdot G - \text{acc.}\mu \cdot M - M'.$$

This reduces the problem to the problem of efficiently computing and updating the commitments. G, M and GT are all commitments to ℓ -sparse vectors, thus can be efficiently computed. The prover can cache the commitments G', M' , and GT' and efficiently update them during accumulation. That is $G'' \leftarrow G' + \alpha G$, $M'' \leftarrow M' + \alpha M$ and $GT'' \leftarrow GT' + \alpha GT$. Additionally, we need to update the accumulation witnesses: $\text{acc}'.\mathbf{m} \leftarrow \text{acc}.\mathbf{m} + \alpha\pi.\mathbf{m}$ and $\text{acc}'.\mathbf{g} \leftarrow \text{acc}.\mathbf{g} + \alpha\pi.\mathbf{g}$. Again because $\pi.\mathbf{g}, \pi.\mathbf{m}$ are sparse this can be done in time $O(\ell)$ independent of $T = |\mathbf{t}|$.

When Π_{LK} is used in composition with another special-sound protocol with a higher degree d , the accumulation is made homogeneous using a $(X + \mu)^{d-2}$ factor when computing the error terms. The contribution to the error terms \mathbf{e}_i ($1 \leq i \leq d-1$) is still a linear function in $\text{acc}.\mathbf{g}$, $\text{acc}.\mathbf{m}$ and $\text{acc}.\mathbf{g} \circ \mathbf{t}$, and thus can be computed homomorphically from commitments to these values.

Special-soundness. We prove special-soundness for the perfect complete version of Π_{LK} , the proof for Π_{LK} is almost identical (but even simpler).

Lemma 5. *The perfect complete version of Π_{LK} is $2(\ell + T)$ -special-sound.*

Proof. We construct an extractor Ext that outputs \mathbf{w} . To show that the witness is valid, we look at the $2(\ell + T)$ transcripts that all have \mathbf{w}, \mathbf{m} as the first message but different $(r^{(j)}, \mathbf{h}^{(j)} \in \mathbb{F}^\ell, \mathbf{g}^{(j)} \in \mathbb{F}^T)$ as the second message. Note that by the pigeonhole principle, there must exist a subset of $S \subseteq [2(\ell + T)]$ transcripts such that $|S| = \ell + T$ and $\mathbf{w}_i + r^{(j)} \neq 0$ for all $i \in [\ell]$ and $j \in S$, and $\mathbf{t}_i + r^{(j)} \neq 0$ for all $i \in [T]$ and $j \in S$. For these transcripts, we have that $\mathbf{h}_i = \frac{1}{\mathbf{w}_i + r^{(j)}}$ and $\mathbf{g}_i = \frac{\mathbf{m}_i}{\mathbf{t}_i + r^{(j)}}$. Define the degree $\ell + T - 1$ polynomial

$$p(X) = \prod_{k=1}^{\ell} (X + \mathbf{w}_k) \cdot \prod_{j=1}^T (X + \mathbf{t}_j) \cdot \left(\sum_{i=1}^{\ell} \frac{1}{X + \mathbf{w}_i} - \sum_{i=1}^T \frac{\mathbf{m}_i}{X + \mathbf{t}_i} \right).$$

If $p(X)$ is the zero polynomial then $\sum_{i=1}^{\ell} \frac{1}{X + \mathbf{w}_i} = \sum_{i=1}^T \frac{\mathbf{m}_i}{X + \mathbf{t}_i}$ and by Lemma 4 $(\mathcal{C}_{\text{LK}}; \mathbf{w}) \in \mathcal{R}_{\text{LK}}$. Since we have $\ell + T$ points $r^{(j)}$ at which $p(r_j) = 0$ we get that $p = 0$ and thus that the extracted witness \mathbf{w} is valid. \square

4.4 Vector-valued lookup

In some applications (e.g., simulating bit operations in circuits), we need to support lookup for a vector, i.e., each table value is a vector of field elements. In this section, we adapt the scheme in Section 4.3 to support vector lookups.

Definition 13. *Consider configuration $\mathcal{C}_{\text{VLK}} := (T, \ell, v \in \mathbb{N}, \mathbf{t})$ where ℓ is the number of lookups, and $\mathbf{t} \in (\mathbb{F}^v)^T$ is a lookup table in which the i th ($1 \leq i \leq T$) entry is*

$$\mathbf{t}_i := (\mathbf{t}_{i,1}, \dots, \mathbf{t}_{i,v}) \in \mathbb{F}^v.$$

A sequence of vectors $\mathbf{w} \in (\mathbb{F}^v)^\ell$ is in relation \mathcal{R}_{VLK} if and only if for all $i \in [\ell]$,

$$\mathbf{w}_i := (\mathbf{w}_{i,1}, \dots, \mathbf{w}_{i,v}) \in \mathbf{t}.$$

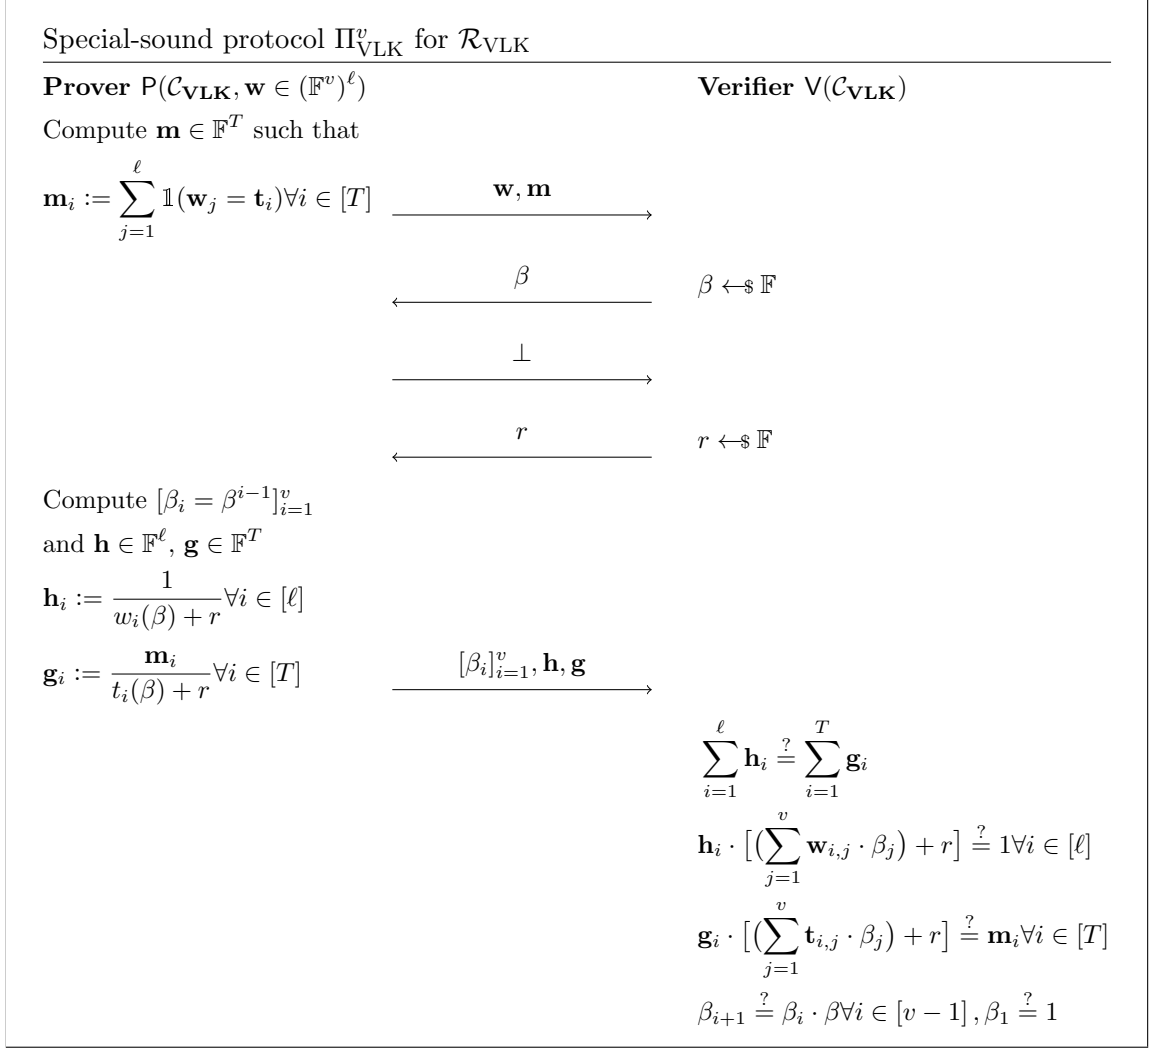
As noted in Section 3.4 of [Hab22], we can extend Lemma 4 and replace Equation 5 with

$$\sum_{i=1}^{\ell} \frac{1}{X + w_i(Y)} = \sum_{i=1}^T \frac{\mathbf{m}_i}{X + t_i(Y)} \quad (6)$$

where the polynomials are defined as

$$w_i(Y) := \sum_{j=1}^v \mathbf{w}_{i,j} \cdot Y^{j-1}, \quad t_i(Y) := \sum_{j=1}^v \mathbf{t}_{i,j} \cdot Y^{j-1},$$

which represent the witness vector $\mathbf{w}_i \in \mathbb{F}^v$ and the table vector $\mathbf{t}_i \in \mathbb{F}^v$. We, therefore, can describe a special-sound protocol for the vector lookup relation as follows.



Achieving perfect completeness. We can use the same trick in Section 4.3 to achieve perfect completeness for Π_{VLK}^v . Namely, the verifier sets $\mathbf{h}_i = 0$ or $\mathbf{g}_i = 0$ when $w_i(\beta) + r = 0$ or $t_i(\beta) + r = 0$ respectively. The verification equations become

$$(w_i(\beta_1, \dots, \beta_v) + r) \cdot (\mathbf{h}_i \cdot (w_i(\beta_1, \dots, \beta_v) + r) - 1) = 0$$

and

$$(t_i(\beta_1, \dots, \beta_v) + r) \cdot (\mathbf{g}_i \cdot (t_i(\beta_1, \dots, \beta_v) + r) - \mathbf{m}_i) = 0,$$

where $w_i(\beta_1, \dots, \beta_v) := \left(\sum_{j=1}^v \mathbf{w}_{i,j} \cdot \beta_j \right)$ and $t_i(\beta_1, \dots, \beta_v) := \left(\sum_{j=1}^v \mathbf{t}_{i,j} \cdot \beta_j \right)$. The degree of the verifier is 5. In practice, the negligible completeness error can likely be ignored without implementing these checks.

Accumulation complexity. Π_{VLK} is a 5-move protocol (i.e. $k = 3$) with the 2nd prover message being empty; the degree of the verifier is 3; the number of non-zero elements in the prover message is at most $(v+3)\ell + v$. To ensure that the accumulation procedure only requires $O(v\ell)$ operations independent of T , we can apply the same trick as in Section 4.3 and compute all error terms from homomorphic commitments to \mathbf{g} and \mathbf{m} . This works because the error terms are a linear function of \mathbf{g} and \mathbf{m} and scalars. This means the contributions of the not necessarily sparse $\text{acc.g}, \text{acc.m}$ to $\mathbf{e}_1, \mathbf{e}_2$ can be computed using the commitment homomorphism.

Special-soundness. We prove that the perfect complete version of Π_{VLK}^v is special-sound.

Lemma 6. *For any $v \in \mathbb{N}$, the perfect complete version of Π_{VLK}^v is $[1 + (v-1) \cdot (\ell + T - 1), 2(\ell + T)]$ -special-sound.*

Proof. We construct an extractor Ext that outputs \mathbf{w} . To show that the witness is valid, we look at the $[1 + (v-1) \cdot (\ell + T - 1), 2(\ell + T)]$ -tree of accepting transcripts. Note that for each depth-1 internal node u that fixes the message $(\mathbf{w}, \mathbf{m}, \beta)$, it has $2(\ell + T)$ different choices of challenge $r^{(j)}$. By the pigeonhole principle, there exists at least $\ell + T$ challenges r such that $t_i(\beta) + r \neq 0$ for all $i \in [T]$ and $w_i(\beta) + r \neq 0$ for all $i \in [\ell]$. Let \mathbf{h}, \mathbf{g} be the last prover message in the corresponding leaf node. Since the transcript is accepting, we have that $\mathbf{h}_i = 1/(w_i(\beta) + r)$ for all $i \in [\ell]$, $\mathbf{g}_i = \mathbf{m}_i/(t_i(\beta) + r)$ for all $i \in [T]$, and $\sum_{i=1}^{\ell} \mathbf{h}_i = \sum_{i=1}^T \mathbf{g}_i$.

Define the bivariate polynomial where the degree of X is $\ell + T - 1$ and the degree of Y is at most $(v-1) \cdot (\ell + T - 1)$,

$$p(X, Y) = \prod_{k=1}^{\ell} (X + w_k(Y)) \cdot \prod_{j=1}^T (X + t_j(Y)) \cdot \left(\sum_{i=1}^{\ell} \frac{1}{X + w_i(Y)} - \sum_{i=1}^T \frac{\mathbf{m}_i}{X + t_i(Y)} \right).$$

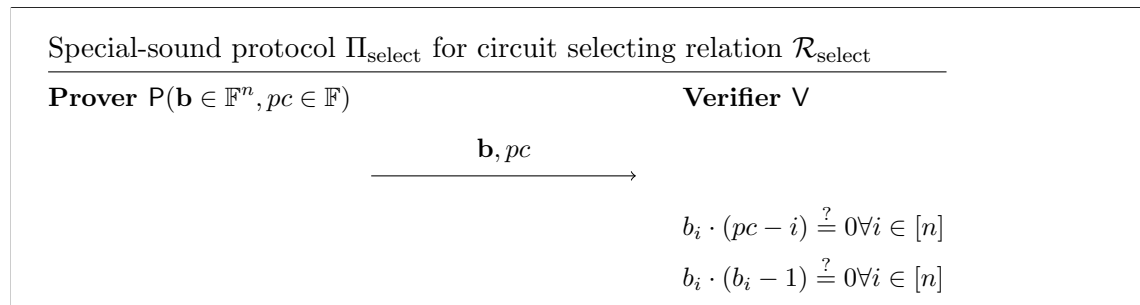
For every depth-1 internal node u , we denote by (r, β) the partial transcript for one of the u 's children whose challenge r satisfies $t_i(\beta) + r \neq 0$ for all $i \in [T]$ and $w_i(\beta) + r \neq 0$ for all $i \in [\ell]$. As argued in the previous paragraph, we observe that $\sum_{i=1}^{\ell} \frac{1}{r + w_i(\beta)} - \sum_{i=1}^T \frac{\mathbf{m}_i}{r + t_i(\beta)} = 0$, hence p evaluates to zero at point (r, β) . Note that there are $(v-1) \cdot (\ell + T - 1) + 1$ depth-1 internal nodes (i.e. $(v-1) \cdot (\ell + T - 1) + 1$ different β s) and each node has $\ell + T$ children (i.e. $\ell + T$ different r) such that p evaluates to zero at point (r, β) . Hence p is the zero polynomial and $\sum_{i=1}^{\ell} \frac{1}{X + w_i(Y)} = \sum_{i=1}^T \frac{\mathbf{m}_i}{X + t_i(Y)}$. Then by the extension of Lemma 4 described in Equation 6, we have $(\mathcal{C}_{\text{VLK}}, \mathbf{w}) \in \mathcal{R}_{\text{VLK}}$ and the extracted witness is valid. \square

4.5 Circuit selection

We provide a sub-protocol for showing that a vector has a single one-bit (and zeros otherwise) at the location of a program counter pc . This is later used to select the appropriate

circuit.

Definition 14. For an integer n the relation $\mathcal{R}_{\text{select}}$ is the set of tuples $(\mathbf{b}, pc) \in \mathbb{F}^n \times \mathbb{F}$ such that $b_i = 0 \forall i \in [n] \setminus \{pc\}$ and if $pc \in [n]$ then $b_{pc} = 1$.



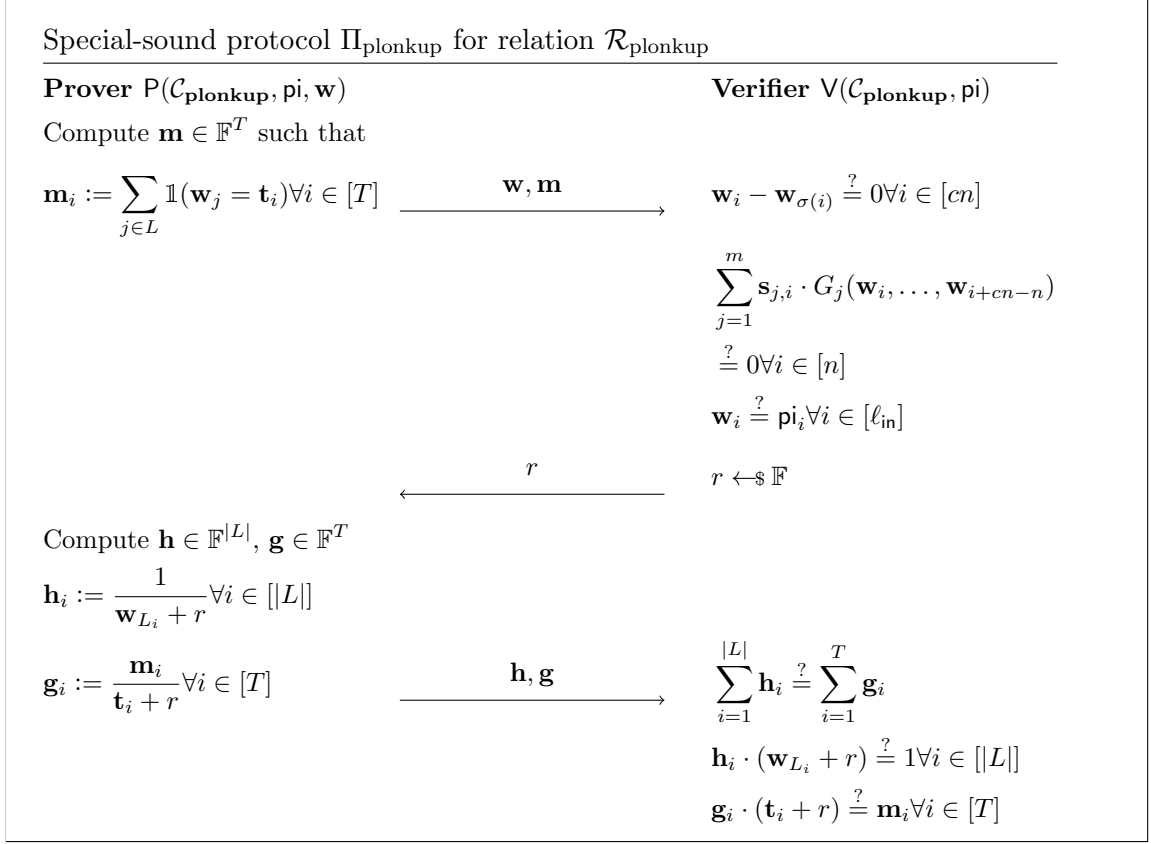
Complexity and security. Π_{select} is a 1-move protocol (i.e. $k = 1$); the degree of the verifier is 2.

5 Special-sound protocols for Plonkup relations

Definition 15. Consider configuration $\mathcal{C}_{\text{plonkup}} := (n, T; \sigma; c, d, [\mathbf{s}_i, G_i]_{i=1}^m; L, \mathbf{t})$ where $\sigma : [cn] \rightarrow [cn]$ is a permutation, $(c, d, [\mathbf{s}_i, G_i]_{i=1}^m)$ are the parameters for the high-degree custom gates, $L \subseteq [cn]$ is the subset of indices for variables that have a lookup gate, $\mathbf{t} \in \mathbb{F}^T$ is the lookup table. The relation $\mathcal{R}_{\text{plonkup}}$ is the set of tuples $(\mathbf{p}_i \in \mathbb{F}^{\ell_{\text{in}}}, \mathbf{w} \in \mathbb{F}^{cn})$ such that

$$\mathbf{w} \in \mathcal{R}_\sigma \wedge \mathbf{w} \in \mathcal{R}_{\text{GATE}} \wedge \mathbf{w}_L \in \mathcal{R}_{\text{LK}} \wedge \mathbf{w}[1..\ell_{\text{in}}] = \mathbf{p}_i.$$

We present the special-sound protocol for the Plonkup relation $\mathcal{R}_{\text{plonkup}}$ below.



Complexity. Π_{plonkup} is a 3-move protocol (i.e. $k = 2$); the degree of the verifier is d ; the number of non-zero elements in the prover message is at most $cn + 3|L|$.

Completeness and security. We need to add the checks described in Section 4.3 to achieve perfect completeness. This changes the verification degree to $\max(d, 3)$. Without these checks, the protocol still has all but negligible completeness.

Lemma 7. Π_{plonkup} is $2(T + |L|)$ -special-sound.

Proof. The protocol is a parallel composition of $\Pi_\sigma, \Pi_{\text{GATE}}$ and Π_{LK} plus a public input check. In Π_σ and Π_{GATE} , the prover simply sends the witness, and the verifier checks it is in the relation. These protocols are thus trivially 1-special-sound, i.e. perfectly sound. The public input relation also trivially holds as the verifier checks $\mathbf{w}_i = \mathbf{pi}_i$ for all $i \in [\ell_{\text{in}}]$. By Lemma 5 Π_{LK} is $2(T + |L|)$ -special-sound. Thus Π_{plonkup} is $2(T + |L|)$ -special-sound. \square

6 Protostar

In this section, we describe PROTOSTAR. PROTOSTAR is built using a special-sound protocol for capturing non-uniform Plonkup circuit computations. In particular, the relation is checking that one of the I circuits is satisfied, where the index of the target circuit is determined by a part of the public input called program counter pc . The non-uniform Plonkup circuit can add arbitrary constraints on input pc . For example, the list of I circuits can be the opcodes supported by EVM, the program counter pc can be computed from the pc' and the register state in the previous step, and the circuit will further check that $\text{opcode}[pc]$ is executed correctly in the current step. For another application, we can consider the I circuits as the predicates of I smart contracts (or transaction types), a user can call one of the smart contracts/transaction types by specifying the index pc , and the cost of proving correct execution is only proportional to the size of an individual smart contract/transaction type rather than the sum of the sizes of the supported smart contracts/transaction types.

For ease of exposition, we assume that the I circuits have the same

- number of gates n ;
- gate arity c ;
- maximum gate degree d ;
- number of gate types m ;
- number of public inputs ℓ_{in} ;
- number of lookup gates ℓ_{k} .

The scheme naturally extends when different branch circuits have different parameters.

Definition 16. Consider configuration $\mathcal{C}_{\text{mplkup}} := (\mathbf{pp} = [n, T, c, d, m, \ell_{\text{in}}, \ell_{\text{k}}]; [\mathcal{C}_i]_{i=1}^I; \mathbf{t})$ where the i th ($1 \leq i \leq I$) branch circuit has configuration $\mathcal{C}_i := (\mathbf{pp}, \sigma_i, [\mathbf{s}_{i,j}, G_{i,j}]_{j=1}^m, L_i)$, and $\mathbf{t} \in \mathbb{F}^T$ is the global lookup table. For a public input $\mathbf{pi} := (pc, \mathbf{pi}') \in \mathbb{F}^{\ell_{\text{in}}}$ where $pc \in [I]$ is a program counter, we say that a instance-witness pair $(\mathbf{pi}, \mathbf{w} \in \mathbb{F}^{cn})$ is in the relation $\mathcal{R}_{\text{mplkup}}$ if and only if $(\mathbf{pi}, \mathbf{w}) \in \mathcal{R}_{\text{plonkup}}$ w.r.t. circuit configuration $(\mathcal{C}_{pc}, \mathbf{t})$.

Protocol $\Pi_{\text{mplkup}} = \langle P(\mathcal{C}_{\text{mplkup}}, \text{pi}, \mathbf{w}), V(\mathcal{C}_{\text{mplkup}}, \text{pi} = (pc \in [I], \text{pi}')) \rangle$:

1. P sends V vector $\mathbf{b} = (0, \dots, 0, b_{pc} = 1, 0, \dots, 0) \in \mathbb{F}^I$.
2. V checks that $b_i \cdot (1 - b_i) \stackrel{?}{=} 0$ and $b_i \cdot (i - pc) \stackrel{?}{=} 0$ for all $i \in [I]$.
3. P computes vector $\mathbf{m} \in \mathbb{F}^T$ such that $\mathbf{m}_i := \sum_{j \in L_{pc}} \mathbb{1}(\mathbf{w}_j = \mathbf{t}_i) \forall i \in [T]$.
4. P sends V vectors \mathbf{w}, \mathbf{m} .
5. V checks that

Permutation check: $\sum_{j=1}^I b_j (\mathbf{w}_i - \mathbf{w}_{\sigma_j(i)}) \stackrel{?}{=} 0$ for all $i \in [cn]$.

Public input check: $\mathbf{w}[1..\ell_{\text{in}}] \stackrel{?}{=} \text{pi}$.

Gate check: for all $i \in [n]$, it holds that

$$\sum_{j=1}^I b_j \cdot \text{GT}_j(\mathbf{s}_{j,1}[i], \dots, \mathbf{s}_{j,m}[i], \mathbf{w}_i, \dots, \mathbf{w}_{i+cn-n}) = 0$$

where $\text{GT}_j(s_1, \dots, s_m, x_1, \dots, x_c) := \sum_{i=1}^m s_i \cdot G_{j,i}(x_1, \dots, x_c)$.

6. V samples and sends P random challenge $r \leftarrow \mathbb{F}$.
7. P computes vectors $\mathbf{h} \in \mathbb{F}^{\ell_{\text{k}}}, \mathbf{g} \in \mathbb{F}^T$ such that

$$\mathbf{h}_i := \frac{1}{\mathbf{w}_{L_{pc}[i]} + r} \forall i \in [\ell_{\text{k}}], \quad \mathbf{g}_i := \frac{\mathbf{m}_i}{\mathbf{t}_i + r} \forall i \in [T].$$

8. V checks that $\sum_{i=1}^{\ell_{\text{k}}} \mathbf{h}_i \stackrel{?}{=} \sum_{i=1}^T \mathbf{g}_i$ and

$$\sum_{j=1}^I b_j \cdot [\mathbf{h}_i \cdot (\mathbf{w}_{L_j[i]} + r)] \stackrel{?}{=} 1 \quad \forall i \in [\ell_{\text{k}}],$$

$$\mathbf{g}_i \cdot (\mathbf{t}_i + r) \stackrel{?}{=} \mathbf{m}_i \quad \forall i \in [T]$$

We present the special-sound protocol Π_{mplkup} for the multi-circuit Plonkup relation.

Remark 4. *By the public input check $\mathbf{w}[1..\ell_{\text{in}}] \stackrel{?}{=} \text{pi}$, we guarantee that $\mathbf{w}[1] = pc$, and the circuit relation can add arbitrary constraints on pc depending on the applications (like the ones described in Section 6).*

Complexity. Π_{mplkup} is a 3-move protocol (i.e. $k = 2$); the degree of the verifier is $d + 1$; the number of non-zero elements in the prover message is at most $cn + 3\ell_{\text{k}} + 1$; the prover message length is $I + cn + 3T$. Hence in the resulting accumulation scheme, the accumulation prover complexity is only $O(n + \ell_{\text{k}})$ that is independent of the table size, and the accumulator size is $O(n + T + I)$ that is independent of the sum of the sizes of the branch circuits. The decider still runs in time $O(I \cdot c \cdot n)$ as it needs to evaluate all circuits at the accumulator.

Special-soundness. We prove the special-soundness property of Π_{mplkup} below.

Lemma 8. Π_{mplkup} is $2(T + \ell_{\text{k}})$ -special-sound.

Proof. The extractor Ext outputs the witness \mathbf{w} sent by the prover. Note that if the verifier checks in step 2 pass, it must be the case that \mathbf{b} is a bool vector with a single non-zero element b_{pc} . Also, note that given $2(T + \ell_{\text{k}})$ accepting transcripts with distinct challenges r , the vector \mathbf{b} won't change. Therefore the sub-transcript after step 2 is essentially a transcript for a Plonkup special-sound protocol Π_{plonkup} with configuration $\mathcal{C}_{\text{plonkup}} := (n, T, c, d, \mathcal{C}_{pc}, \mathbf{t})$. By Lemma 7, it holds that Π_{mplkup} is $2(T + \ell_{\text{k}})$ -special-sound. \square

We will now use Π_{mplkup} and our compiler described in Theorem 3 to design **PROTOSTAR**. Before that, we must address an efficiency issue when combining the high-degree gate and sparse lookup protocols with the generic transform **CV** in Section 3.5.

Error separation between high-degree checks and low-degree checks. In some scenarios, the set of equations that a verifier checks can be with heterogeneous degrees. For example, in Section 3.5, the transformed protocol $\text{CV}[\Pi_{\text{sps}}]$ has one degree- $(d + 2)$ checks and $O(\sqrt{\ell})$ degree-2 checks. In this case, we can partition the checks into two sets whereas one with high-degree equations and the other with low-degree equations. We then assign separate sets of error terms for the two sets of equations. The first set of error terms can use the trivial identity commitment (as shown in Section 3.5) and thus reduces the number of group operations. The rest of the error term vectors can be committed using group elements. Since the degree of the second set is low, the number of error commitments will be a small constant (e.g. 2), thus the number of additional group operations performed in the recursive circuit will be small. We refer to Appendix A for more details.

Efficient accumulation of $\text{CV}[\Pi_{\text{mplkup}}]$. $\text{CV}[\Pi_{\text{GATE}}]$ reduces the number of degree- d verification checks in Π_{GATE} from n to 1, with the tradeoff of $O(\sqrt{n})$ additional degree-2 checks. In the resulting accumulation scheme, the error terms for high-degree gates are, thus, only of length 1. This enables using the trivial identity commitment for these error terms and thus reduces the number of group operations by the accumulation verifier. Unfortunately, applying **CV** to mplkup seems to have a major tradeoff. The number of verification checks is $n + \ell_{\text{k}} + T + c \cdot n$. This requires using a) $\text{CV}[\text{mplkup}]$ and b) is not

composable with the sparseness optimizations for lookup described in Sections 4.3 and 4.4. These optimizations make the prover computation independent of T .

Fortunately, a closer look at the verification of `mplkup` reveals that only n of these verification checks are of high degree d , namely the checks in Π_{GATE} . The other checks are of degree 2 or lower. With a slight abuse of notation, we can define $\text{CV}[\Pi_{\text{mplkup}}]$ as applying the generic transform `CV` only to the Π_{GATE} part of Π_{mplkup} . This means that there are $d + 1$ cross error vectors (each of length 1) for the degree $d + 2$ check in $\text{CV}[\Pi_{\text{GATE}}]$; and 1 cross error vector of length $T + \ell_{\text{lk}} + cn + O(\sqrt{n})$ for the rest checks—namely the low-degree checks in Π_{mplkup} and the $O(\sqrt{n})$ degree-2 checks in $\text{CV}[\Pi_{\text{GATE}}]$. By leveraging the error separation technique described above, we can use the identity function to commit to the field elements and a vector commitment to commit to the long error term. We can again leverage homomorphism as described in Section 4.3 to make the prover independent of T .

Corollary 1 (PROTOSTAR protocol). *Consider the configuration*

$$C_{\text{mplkup}} := (n, T, c, d, m, \ell_{\text{in}}, \ell_{\text{lk}}; [\mathcal{C}_i]_{i=1}^I; \mathbf{t}).$$

Given a binding homomorphic commitment scheme $\text{cm} = (\text{Setup}, \text{Commit})$, and under the Fiat-Shamir Heuristic (Definition 9) for a hash function \mathbf{H} , there exists an IVC scheme PROTOSTAR for $\mathcal{R}_{\text{mplkup}}$ relations with the following efficiencies for $m = 1$ (i.e. each circuit has a single degree- d gate type), public input length $\ell_{\text{in}} = 1$: (we omit cost terms that are negligible compared to the dominant parts)

$\text{P}_{\text{PROTOSTAR}}$ <i>native</i>	$\text{P}_{\text{PROTOSTAR}}$ <i>recursive</i>	$\text{V}_{\text{PROTOSTAR}}$	$ \pi_{\text{PROTOSTAR}} $
$O(\mathbf{w} + \ell_{\text{lk}})\mathbb{G}$ $L'(\mathcal{C}_i, d + 2) + 2\ell_{\text{lk}}\mathbb{F}$	$3\mathbb{G}$ $d + 4\mathbb{F}$ $d + O(1)H + 1H_{\text{in}}$	$O(c \cdot n + T + \ell_{\text{lk}})\mathbb{G}$ $n + \sum_{i=1}^I \mathcal{C}_i + T + \ell_{\text{lk}}\mathbb{F}$	$O(c \cdot n + T + \ell_{\text{lk}})$

Here $|\mathbf{w}| \leq cn$ is the number of non-zero entries in the witness, $\sum_{i=1}^I \mathcal{C}_i$ is the cost of evaluating all circuits on some random input, and $L'(\mathcal{C}_i, d)$ is the cost of computing the coefficients of the polynomial $e(X)$ defined in Equation 4.⁵ H_{in} is the cost of hashing the public input and the (constant-sized) accumulator instance.

Proof. Let $\text{SPS} - \text{IVC}[\Pi] = \text{IVC}[\text{acc}[\text{FS}[\text{cm}[\text{CV}[\Pi]]]]]$ be the transformation from a special-sound protocol to an IVC-scheme described by Theorem 3 (including `CV`). Then given a commitment scheme `cm` by that theorem $\text{PROTOSTAR} = \text{SPS} - \text{IVC}[\Pi_{\text{mplkup}}]$ is an IVC scheme for predicates expressed in $\mathcal{R}_{\text{mplkup}}$. We apply Theorem 3 to get the efficiencies in the table above.

⁵As noted in Theorem 3, $L'(\mathcal{C}_i, d + 2)$ is bounded by $O(nd \log^2(d))$.

Security: Since $\text{CV}[\Pi_{\text{GATE}}]$ is only applied to Π_{GATE} which has perfect soundness, by Lemma 8 and Lemma 3, the NARK scheme $\text{FS}[\text{cm}[\text{CV}[\Pi]]]$ for $\mathcal{R}_{\text{mplkup}}$ has knowledge soundness with knowledge error $(Q + 1) \cdot \frac{n+2(T+\ell_{\text{ik}})}{|\mathbb{F}|} + \text{negl}(\lambda)$, where Q is the number of RO queries by the adversary. Using Theorem 2 and Corollary 2 we can construct an accumulation scheme for the NARK scheme $\text{FS}[\text{cm}[\text{CV}[\Pi]]]$. The accumulation scheme has negligible knowledge error as $d = \text{poly}(\lambda)$. Therefore, under the Fiat-Shamir heuristic and by Theorem 1, $\text{SPS} - \text{IVC}[\Pi]$ is a secure IVC scheme. \square

Acknowledgments. We would like to thank Ariel Gabizon and Liam Eagen for the inspiring discussions about optimizing the generic transformation protocol in Section 3.5. We'd like to thank Zachary Williamson, Sean Bowe, Srinath Setty and Nicholas Mohnblatt for pointing out typos and minor mistakes.

References

- [AFGHO10] Masayuki Abe, Georg Fuchsbauer, Jens Groth, Kristiyan Haralambiev, and Miyako Ohkubo. “Structure-Preserving Signatures and Commitments to Group Elements”. In: *CRYPTO 2010*. Ed. by Tal Rabin. Vol. 6223. LNCS. Springer, Heidelberg, Aug. 2010, pp. 209–236. DOI: 10.1007/978-3-642-14623-7_12.
- [AFK22] Thomas Attema, Serge Fehr, and Michael Klooß. “Fiat-Shamir Transformation of Multi-round Interactive Proofs”. In: *TCC 2022, Part I*. Ed. by Eike Kiltz and Vinod Vaikuntanathan. Vol. 13747. LNCS. Springer, Heidelberg, Nov. 2022, pp. 113–142. DOI: 10.1007/978-3-031-22318-1_5.
- [BBBF18] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. “Verifiable Delay Functions”. In: *CRYPTO 2018, Part I*. Ed. by Hovav Shacham and Alexandra Boldyreva. Vol. 10991. LNCS. Springer, Heidelberg, Aug. 2018, pp. 757–788. DOI: 10.1007/978-3-319-96884-1_25.
- [BCCT13] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. “Recursive composition and bootstrapping for SNARKS and proof-carrying data”. In: *45th ACM STOC*. Ed. by Dan Boneh, Tim Roughgarden, and Joan Feigenbaum. ACM Press, June 2013, pp. 111–120. DOI: 10.1145/2488608.2488623.
- [BCLMS21] Benedikt Bünz, Alessandro Chiesa, William Lin, Pratyush Mishra, and Nicholas Spooner. “Proof-Carrying Data Without Succinct Arguments”. In: *CRYPTO 2021, Part I*. Ed. by Tal Malkin and Chris Peikert. Vol. 12825. LNCS. Virtual Event: Springer, Heidelberg, Aug. 2021, pp. 681–710. DOI: 10.1007/978-3-030-84242-0_24.
- [BCMS20] Benedikt Bünz, Alessandro Chiesa, Pratyush Mishra, and Nicholas Spooner. “Recursive Proof Composition from Accumulation Schemes”. In: *TCC 2020, Part II*. Ed. by Rafael Pass and Krzysztof Pietrzak. Vol. 12551. LNCS. Springer, Heidelberg, Nov. 2020, pp. 1–18. DOI: 10.1007/978-3-030-64378-2_1.
- [BCTV14] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. “Scalable Zero Knowledge via Cycles of Elliptic Curves”. In: *CRYPTO 2014, Part II*. Ed. by Juan A. Garay and Rosario Gennaro. Vol. 8617. LNCS. Springer, Heidelberg, Aug. 2014, pp. 276–294. DOI: 10.1007/978-3-662-44381-1_16.
- [BGH19] Sean Bowe, Jack Grigg, and Daira Hopwood. *Halo: Recursive Proof Composition without a Trusted Setup*. Cryptology ePrint Archive, Report 2019/1021. <https://eprint.iacr.org/2019/1021>. 2019.

- [BMRS20] Joseph Boneau, Izaak Meckler, Vanishree Rao, and Evan Shapiro. *Coda: Decentralized Cryptocurrency at Scale*. Cryptology ePrint Archive, Report 2020/352. <https://eprint.iacr.org/2020/352>. 2020.
- [But22] Vitalik Buterin. *The different types of ZK EVM*. <https://vitalik.ca/general/2022/08/04/zkevm.html>. Accessed: 2023-04-27. 2022.
- [CBBZ22] Binyi Chen, Benedikt Bünz, Dan Boneh, and Zhenfei Zhang. *HyperPlonk: Plonk with Linear-Time Prover and High-Degree Custom Gates*. Cryptology ePrint Archive, Report 2022/1355. <https://eprint.iacr.org/2022/1355>. 2022.
- [CT10] Alessandro Chiesa and Eran Tromer. “Proof-Carrying Data and Hearsay Arguments from Signature Cards”. In: *ICS 2010*. Ed. by Andrew Chi-Chih Yao. Tsinghua University Press, Jan. 2010, pp. 310–331.
- [CTV15] Alessandro Chiesa, Eran Tromer, and Madars Virza. “Cluster Computing in Zero Knowledge”. In: *EUROCRYPT 2015, Part II*. Ed. by Elisabeth Oswald and Marc Fischlin. Vol. 9057. LNCS. Springer, Heidelberg, Apr. 2015, pp. 371–403. DOI: 10.1007/978-3-662-46803-6_13.
- [EFG22] Liam Eagen, Dario Fiore, and Ariel Gabizon. *cq: Cached quotients for fast lookups*. Cryptology ePrint Archive, Report 2022/1763. <https://eprint.iacr.org/2022/1763>. 2022.
- [GW20] Ariel Gabizon and Zachary J. Williamson. *plookup: A simplified polynomial protocol for lookup tables*. Cryptology ePrint Archive, Report 2020/315. <https://eprint.iacr.org/2020/315>. 2020.
- [GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. *PLONK: Permutations over Lagrange-bases for Oecumenical Noninteractive arguments of Knowledge*. Cryptology ePrint Archive, Report 2019/953. <https://eprint.iacr.org/2019/953>. 2019.
- [Hab22] Ulrich Haböck. *Multivariate lookups based on logarithmic derivatives*. Cryptology ePrint Archive, Report 2022/1530. <https://eprint.iacr.org/2022/1530>. 2022.
- [KB20] Assimakis Kattis and Joseph Boneau. *Proof of Necessary Work: Succinct State Verification with Fairness Guarantees*. Cryptology ePrint Archive, Report 2020/190. <https://eprint.iacr.org/2020/190>. 2020.
- [KMT22] Dmitry Khovratovich, Mary Maller, and Pratyush Ranjan Tiwari. *Min-Root: Candidate Sequential Function for Ethereum VDF*. Cryptology ePrint Archive, Report 2022/1626. <https://eprint.iacr.org/2022/1626>. 2022.

- [KS22] Abhiram Kothapalli and Srinath Setty. *SuperNova: Proving universal machine executions without universal circuits*. Cryptology ePrint Archive, Report 2022/1758. <https://eprint.iacr.org/2022/1758>. 2022.
- [KS23] Abhiram Kothapalli and Srinath Setty. “HyperNova: Recursive arguments for customizable constraint systems”. In: *Cryptology ePrint Archive* (2023).
- [KST22] Abhiram Kothapalli, Srinath Setty, and Ioanna Tzialla. “Nova: Recursive Zero-Knowledge Arguments from Folding Schemes”. In: *CRYPTO 2022, Part IV*. Ed. by Yevgeniy Dodis and Thomas Shrimpton. Vol. 13510. LNCS. Springer, Heidelberg, Aug. 2022, pp. 359–388. DOI: 10.1007/978-3-031-15985-5_13.
- [LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. “Algebraic methods for interactive proof systems”. In: *Journal of the ACM (JACM)* 39.4 (1992), pp. 859–868.
- [Moh23] Nicholas Mohnblatt. *Sangria: A Folding Scheme for PLONK*. https://github.com/geometryresearch/technical_notes/blob/main/sangria_folding_plonk.pdf. Accessed: 2023-04-27. 2023.
- [NT16] Assa Naveh and Eran Tromer. “PhotoProof: Cryptographic Image Authentication for Any Set of Permissible Transformations”. In: *2016 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2016, pp. 255–271. DOI: 10.1109/SP.2016.23.
- [Ped92] Torben P. Pedersen. “Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing”. In: *CRYPTO’91*. Ed. by Joan Feigenbaum. Vol. 576. LNCS. Springer, Heidelberg, Aug. 1992, pp. 129–140. DOI: 10.1007/3-540-46766-1_9.
- [PK22] Jim Posen and Assimakis A. Kattis. *Caulk+: Table-independent lookup arguments*. Cryptology ePrint Archive, Report 2022/957. <https://eprint.iacr.org/2022/957>. 2022.
- [SAGL18] Srinath Setty, Sebastian Angel, Trinabh Gupta, and Jonathan Lee. “Proving the correct execution of concurrent services in zero-knowledge”. In: *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. 2018, pp. 339–356.
- [STW23] Srinath Setty, Justin Thaler, and Riad Wahby. “Customizable constraint systems for succinct arguments”. In: *Cryptology ePrint Archive* (2023).
- [Val08] Paul Valiant. “Incrementally Verifiable Computation or Proofs of Knowledge Imply Time/Space Efficiency”. In: *TCC 2008*. Ed. by Ran Canetti. Vol. 4948. LNCS. Springer, Heidelberg, Mar. 2008, pp. 1–18. DOI: 10.1007/978-3-540-78524-8_1.

- [Wik21] Douglas Wikström. *Special Soundness in the Random Oracle Model*. Cryptology ePrint Archive, Report 2021/1265. <https://eprint.iacr.org/2021/1265>. 2021.
- [XCZBFKC22] Alex Luoyuan Xiong, Binyi Chen, Zhenfei Zhang, Benedikt Bünz, Ben Fisch, Fernando Krell, and Philippe Camacho. *VERI-ZEXE: Decentralized Private Computation with Universal Setup*. Cryptology ePrint Archive, Report 2022/802. <https://eprint.iacr.org/2022/802>. 2022.
- [ZBKMNS22] Arantxa Zapico, Vitalik Buterin, Dmitry Khovratovich, Mary Maller, Anca Nitulescu, and Mark Simkin. “Caulk: Lookup Arguments in Sublinear Time”. In: *ACM CCS 2022*. Ed. by Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi. ACM Press, Nov. 2022, pp. 3121–3134. DOI: 10.1145/3548606.3560646.
- [ZGKMR22] Arantxa Zapico, Ariel Gabizon, Dmitry Khovratovich, Mary Maller, and Carla Ràfols. *Baloo: Nearly Optimal Lookup Arguments*. Cryptology ePrint Archive, Report 2022/1565. <https://eprint.iacr.org/2022/1565>. 2022.

A Accumulation Scheme for high/low degree verifier

We describe a modification for the accumulation scheme in Section 3.4 that can be useful if V_{sps} has both a single high-degree verification check and multiple low-degree checks. E.g. assume that $V_{\text{sps}} = V_{\text{sps},1} || V_{\text{sps},2}$ where $V_{\text{sps},1} : (\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}) \rightarrow \mathbb{F}$ is degree d and $V_{\text{sps},2} : (\mathbf{pi}, [\mathbf{m}_i]_{i=1}^k, [r_i]_{i=1}^{k-1}) \rightarrow \mathbb{F}^\ell$ is degree 2.

For simplicity, we assume that $V_{\text{sps},1}$ maps to a single field element and that $V_{\text{sps},2}$ is degree 2, but this naturally extends to more arbitrary degrees, sizes and more components.

The accumulation scheme $\text{acc}_{\text{HL}} = (P_{\text{acc,HL}}, V_{\text{acc,HL}})$ for $\text{FS}[V_{\text{sps},1} || V_{\text{sps},2}]$ is essentially a parallel composition of the accumulation presented Section 3.4 applied to $V_{\text{sps},1}$ and $V_{\text{sps},2}$. Concretely there are the following modifications to the scheme from Section 3.4:

- The prover computes error terms separately for $V_{\text{sps},1}$ and $V_{\text{sps},2}$. This means there are $d - 1$ constant size error terms and 1 error term vector of length ℓ .
- The prover uses the identity function to commit to the d error terms and a homomorphic vector commitment $\text{cm} = (\text{Setup}, \text{Commit})$ to commit to the single length ℓ error term.
- The accumulator stores two error terms, one for each verifier. One is a field element e , and the other is a commitment E to a length ℓ vector.
- The accumulation verifier checks the correct accumulation for each error term separately, thus performing $d - 1$ field operations and 1 homomorphic commitment scalar multiplication.

Complexity and security. The scheme has the following complexity:

- The *accumulation prover*

- asks $k - 1$ queries to ρ_{NARK} with constant-sized inputs and 1 query to ρ_{acc} with input size $d + O(1)$;
- computes the coefficients of

$$e(X) = \sum_{j=0}^d (\mu + X)^{d-j} f_j^{\text{V}_{\text{sps},1}}(X \cdot \text{pi} + \text{acc.pi}, [X \cdot \mathbf{m}_i + \text{acc.m}_i]_{i=1}^k, [X \cdot r_i + \text{acc.r}_i]_{i=1}^{k-1}) \in \mathbb{F}[X]$$

and computes $\mathbf{e} \in \mathbb{F}^\ell$, which are the coefficients of X in the polynomials

$$\sum_{j=0}^2 (\mu + X)^{2-j} f_j^{\text{V}_{\text{sps},2}}(X \cdot \text{pi} + \text{acc.pi}, [X \cdot \mathbf{m}_i + \text{acc.m}_i]_{i=1}^k, [X \cdot r_i + \text{acc.r}_i]_{i=1}^{k-1}) \in (\mathbb{F}[X])^\ell$$

- commits to \mathbf{e} using ℓ \mathbb{G} ops.
- performs $|R| + |M^*| + 2$ \mathbb{F} -ops to combine $(\mu, \text{pi}, [r_i]_{i=1}^{k-1}, [\mathbf{m}_i]_{i=1}^k)$ (where $|R|$ is the number of challenges and $|M^*|$ is the number of non-zero elements in prover messages);
- performs k \mathbb{G} -ops to combine $[C_i]_{i=1}^k$;
- performs 1 \mathbb{G} -op to add $E = \text{Commit}(\text{ck}, \mathbf{e})$ to acc.E .
- The *accumulation verifier* performs
 - asks $k - 1$ queries to ρ_{NARK} and 1 query to ρ_{acc} ;
 - $|R| + 2$ \mathbb{F} -ops to combine $(\mu, \text{pi}, [r_i]_{i=1}^{k-1})$;
 - k \mathbb{G} -ops to combine $[C_i]_{i=1}^k$;
 - $d - 1$ \mathbb{F} -ops to add $[e_j]_{j=1}^{d-1}$ onto acc.e .
 - 1 \mathbb{G} -op to add $E = \text{Commit}(\text{ck}, \mathbf{e})$ to acc.E
- The *decider*
 - computes $C_i = \text{Commit}(\text{ck}, \mathbf{m}_i)$ for $i \in [k]$ and $E = \text{Commit}(\text{ck}, \mathbf{e})$, with total complexity around $|M| + \ell$ \mathbb{G} -ops.
 - evaluate $[f_i^{\text{V}_{\text{sps},1}}]_{i=0}^d$ and $[f_i^{\text{V}_{\text{sps},2}}]_{i=0}^2$ to verify e and \mathbf{e} .

Corollary 2 (Hi-LowAccumulation). *Let $(\text{P}_{\text{NARK,HL}}, \text{V}_{\text{NARK,HL}}) = \text{FS}[\text{V}_{\text{sps},1} || \text{V}_{\text{sps},2}]$ be an RO-NARK as defined above. Let cm be a binding, homomorphic commitment scheme and ρ_{acc} be a random oracle. The accumulation scheme acc_{HL} for V_{NARK} satisfies perfect completeness and has knowledge-error $(Q + 1) \frac{d+4}{|\mathbb{F}|} + \text{negl}(\lambda)$, as defined in Definition 8.*

Proof sketch: Perfect completeness follows immediately from Theorem 2. For knowledge-soundness, consider that acc_{HL} is a parallel composition of two accumulation schemes applied to a high-degree and a low-degree verifier. Given an adversary that can break the knowledge soundness of acc_{HL} , we can construct an adversary that can break the knowledge soundness of either the high or the low-degree accumulation. By union bound, this leads to a knowledge error of $(Q + 1) \frac{d+4}{|\mathbb{F}|} + \text{negl}(\lambda)$.

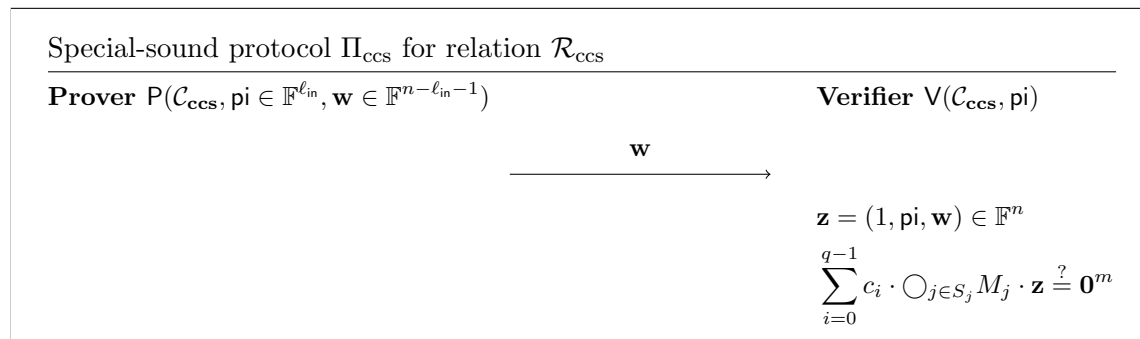
B Protostar for CCS

Recently, Setty, Thaler and Wahby introduced Customizable Constraint System (CCS), a new characterization of NP that is a generalization of R1CS[STW23]. It enables the use of high-degree gates while not requiring permutation arguments. It is also powerful enough to capture both R1CS and Plonk constraints. As described in the introduction, HyperNova builds an accumulation scheme and, thus, IVC for CCS. However, HyperNova does not natively support non-uniform circuits and both the recursion cost, as well as the cost for lookups is more expensive than PROTOSTAR, which is built for mplkup. However, mplkup still requires a permutation argument and so-called copy-constraints. We show that our general compiler is powerful enough to port the benefits of PROTOSTAR directly to CCS. The starting point is the trivial special-sound protocol for CCS:

Definition 17 ([STW23]). *Given public parameters $m, n, N, \ell_{\text{in}}, t, q, d \in \mathbb{N}$ where $n > \ell$, let $M_1, \dots, M_t \in \mathbb{F}^{m \times n}$ be matrices with at most N total non-zero entries. Let S_1, \dots, S_q be multisets over domain $[t]$ and each multiset has cardinality at most d . Let $c_1, \dots, c_q \in \mathbb{F}$ be constants. A tuple $(\mathbf{pi}, \mathbf{w}) \in \mathbb{F}^{\ell_{\text{in}}} \times \mathbb{F}^{n - \ell_{\text{in}} - 1}$ is in the relation \mathcal{R}_{CCS} if and only if*

$$\sum_{i=0}^{q-1} c_i \cdot \bigcirc_{j \in S_j} M_j \cdot \mathbf{z} = \mathbf{0}^m,$$

where $z = (1, \mathbf{pi}, \mathbf{w})$. \bigcirc denotes the Hadamard product between vectors.



Complexity. Π_{ccs} is a 1-move protocol (i.e. $k = 1$) with verifier degree d .

Next, we present the special-sound protocol $\Pi_{\text{mccs}+}$ for the extended CCS relation that has multi-circuits and lookup support. Compared to Π_{implkup} in Section 6, the protocol $\Pi_{\text{mccs}+}$ removes the need of a permutation check.

Definition 18. Consider configuration $\mathcal{C}_{\text{mccs}+} := (\text{pp} = [m, n, N, \ell_{\text{in}}, t, q, d, T, \ell_{\text{lk}}]; [\mathcal{C}_i]_{i=1}^I; \mathbf{t})$ where the i th ($1 \leq i \leq I$) branch circuit has configuration $\mathcal{C}_i := (\text{pp}, [M_{j,i}]_{j=1}^t, [S_{j,i}, c_{j,i}]_{j=1}^q, L_i)$, and $\mathbf{t} \in \mathbb{F}^T$ is the global lookup table. For a public input $\text{pi} := (pc, \text{pi}') \in \mathbb{F}^{\ell_{\text{in}}}$ where $pc \in [I]$ is a program counter, we say that a instance-witness pair $(\text{pi}, \mathbf{w}) \in \mathbb{F}^{n-1}$ is in the relation $\mathcal{R}_{\text{mccs}+}$ if and only if $(\text{pi}, \mathbf{w}) \in \mathcal{R}_{\text{ccs}}$ w.r.t. circuit configuration \mathcal{C}_{pc} and $\mathbf{w}_{L_{pc}} \subseteq \mathbf{t}$.

Protocol $\Pi_{\text{mccs}+} = \langle \text{P}(\mathcal{C}_{\text{mccs}+}, \text{pi}, \mathbf{w}), \text{V}(\mathcal{C}_{\text{mccs}+}, \text{pi} = (pc \in [I], \text{pi}')) \rangle$:

1. P sends V vector $\mathbf{b} = (0, \dots, 0, b_{pc} = 1, 0, \dots, 0) \in \mathbb{F}^I$.
2. V checks that $b_i \cdot (1 - b_i) \stackrel{?}{=} 0$ and $b_i \cdot (i - pc) \stackrel{?}{=} 0$ for all $i \in [I]$.
3. P computes vector $\mathbf{m} \in \mathbb{F}^T$ such that $\mathbf{m}_i := \sum_{j \in L_{pc}} \mathbb{1}(\mathbf{w}_j = \mathbf{t}_i) \forall i \in [T]$.
4. P sends V vectors \mathbf{w}, \mathbf{m} .
5. V computes $\mathbf{z} = (1, \text{pi}, \mathbf{w}) \in \mathbb{F}^n$ and checks that

$$\sum_{k=1}^I b_k \cdot \sum_{i=1}^q c_{i,k} \cdot \bigcirc_{j \in S_{i,k}} M_{j,k} \cdot \mathbf{z} = \mathbf{0}^m.$$

6. V samples and sends P random challenge $r \leftarrow \mathbb{F}$.
7. P computes vectors $\mathbf{h} \in \mathbb{F}^{\ell_{\text{lk}}}, \mathbf{g} \in \mathbb{F}^T$ such that

$$\mathbf{h}_i := \frac{1}{\mathbf{w}_{L_{pc}[i]} + r} \forall i \in [\ell_{\text{lk}}], \quad \mathbf{g}_i := \frac{\mathbf{m}_i}{\mathbf{t}_i + r} \forall i \in [T].$$

8. V checks that $\sum_{i=1}^{\ell_{\text{lk}}} \mathbf{h}_i \stackrel{?}{=} \sum_{i=1}^T \mathbf{g}_i$ and

$$\sum_{j=1}^I b_j \cdot \left[\mathbf{h}_i \cdot (\mathbf{w}_{L_j[i]} + r) \right] \stackrel{?}{=} 1 \quad \forall i \in [\ell_{\text{lk}}],$$

$$\mathbf{g}_i \cdot (\mathbf{t}_i + r) \stackrel{?}{=} \mathbf{m}_i \quad \forall i \in [T].$$

Complexity. $\Pi_{\text{mccs}+}$ is a 3-move protocol (i.e. $k = 2$); the degree of the verifier is $d + 1$; the number of non-zero elements in the prover message is at most $n + 3\ell_{\text{lk}}$; the prover

message length is at most $I + n + 3T$. Hence in the resulting accumulation scheme, the accumulation prover complexity is only $O(n + \ell_{\mathbb{k}})$ that is independent of the table size, and the accumulator size is $O(n + T + I)$ that is independent of the sum of the sizes of the branch circuits. We detail the efficiency of the resulting IVC scheme in the table below. The efficiency is almost identical to the PROTOSTAR scheme for $\mathcal{R}_{\text{mplkup}}$. However, the cost of computing the error terms $L'(\text{ccs}_i, d + 2)$ now depends on the i th CCS instance.

$P_{\text{PROTOSTAR, mccs+ native}}$	$P_{\text{PROTOSTAR, mccs+ recursive}}$	$V_{\text{PROTOSTAR, mccs+}}$	$ \pi_{\text{PROTOSTAR, mccs+}} $
$O(\mathbf{w} + \ell_{\mathbb{k}})\mathbb{G}$ $L'(\text{ccs}_i, d + 2) + 2\ell_{\mathbb{k}}\mathbb{F}$	$3\mathbb{G}$ $d + 4\mathbb{F}$ $d + O(1)H + 1H_{\text{in}}$	$O(n + T + \ell_{\mathbb{k}})\mathbb{G}$ $n + I \cdot N + T + \ell_{\mathbb{k}}\mathbb{F}$	$O(n + T + \ell_{\mathbb{k}})$